arXiv:1406.1131v1 [math.PR] 4 Jun 2014

The Metric Coalescent

Daniel Lanoue

June 5, 2014

Abstract

The Metric Coalescent (MC) is a measure-valued Markov Process generalizing the classical Kingman Coalescent. We show how the MC arises naturally from a discrete agent based model (Compulsive Gambler) of social dynamics and prove an existence and uniqueness theorem extending the MC to the space of all Borel probability measures on any locally compact Polish space.

Contents

1 Introduction **2**
- 1.1 Setup . . . 3
- 1.2 The Metric Coalescent . . . 3
- 1.3 Extension to $P(S)$ . . . 4
- 1.4 Overview . . . 5

2 The Token Process **6**
- 2.1 Construction . . . 6
- 2.2 The Empirical Measures . . . 7
- 2.3 Exchangeable Partitions, Measure-Valued Markov Processes . . . 8
- 2.4 Two Types of Exchangeability . . . 9
  - 2.4.1 Symmetric Exchangeability . . . 9
  - 2.4.2 Asymmetric Exchangeability . . . 11
- 2.5 Proof of 2.1 . . . 14
  - 2.5.1 Kingman's Paintbox Theorem . . . 15
  - 2.5.2 Defining μ_t . . . 15
  - 2.5.3 Convergence . . . 17
- 2.6 Support and Total Variation Distance . . . 18

3 The Markov Property **20**
- 3.1 Markov Condition for μ_t^N . . . 20
- 3.2 Proof of 3.1 . . . 23

4 Finite Support and Kingman's Coalescent **25**
- 4.1 Kingman's Coalescent . . . 26
- 4.2 Coming Down From Infinity . . . 26

5 The Metric Coalescent **28**
5.1 The Partial Order $\prec$ 28
5.2 Size Biased Ordering 29
5.2.1 Definition of Size Biased Ordering 29
5.3 Proof of 5.1 30

6 Martingales of the Metric coalescent **30**
6.1 Process Dynamics of the Metric Coalescent 31
6.2 Second Moment Calculation 32
6.3 Proof of 6.1 35

7 Feller Continuity **36**
7.1 Convergence of Random Measures 37
7.2 Proof of 7.2 38
7.3 The Coupling 40
7.4 Proof of 7.4 41
7.5 Moment Calculations for $\mu_t(f)$ 49

8 Continuity and Uniqueness **50**
8.1 Proof of 8.1 50
8.2 Uniqueness of the Extension 52

9 Some Motivating Examples **53**
9.1 The Rate $\phi(x)$ Near Zero 53
9.1.1 An Example 53
9.2 Finite Support for Positive Times 54
9.2.1 An Example 55

10 Further Directions **56**
10.1 Coming Down From Infinity 56
10.2 Time Reversal 57

1 Introduction

Introduced in [3] the Compulsive Gambler (CG) process is an Finite Markov Information Exchange (FMIE) process modelling a finite set of agents meeting pairwise randomly and playing a fair winner-take-all game. The precise set-up of the model in the FMIE framework is two leveled. The first level consists of the agents and their meeting model, described by the finite set Agents and a non-negative array (ν_{ij}) of meeting rates indexed by unordered pairs $\{i, j\}$ of agents. A pair of agents i, j then meet at the times an independent Poisson process of rate ν_{ij}.

The second level in the FMIE framework is the information exchange model, which describes the state $X_i(t)$ of each agent i at time t and the (deterministic or random) update rule for the states of two agents i and j upon meeting. In the CG process, when two agents with non-zero wealth a and b meet they play a fair game in which one agent acquires the combined amount $a + b$ and the other is left with nothing (thus with probabilities $\frac{a}{a+b}$ and $\frac{b}{a+b}$ respectively). Normalizing the total wealth present amongst the agents, the state of the CG process can be viewed at each time as a probability measure over Agents.

Several directions for research concerning the CG process are suggested in [3] and this paper is devoted to a detailed study of one direction, the extension to continuous state space. We introduce the Metric Coalescent (MC), a measure valued Markov process defined over any metric space (S, d), a generalization of the CG process given by adding some geometry to the meeting model between agents. Viewing each agent as occupying a location in S, their meeting rates are then determined by the distances between them.

Our interest in the Metric Coalescent comes from the wide variety of techniques that can be exploited in its study. As a natural generalization of an FMIE process, we can apply tools from the study of finite Markov chains (or more generally Interacting Particle Systems) and in particular study the process dynamics of certain martingales associated to the process. As the MC is constructed from (S, d) in terms of an exchangeable random distance array (as in [19]), tools from the theory of exchangeability such de Finetti's Theorem and Kingman's Paintbox Theorem are applied. Finally, viewed through a symmetry breaking duality with an exchangeable partition, we make use of a comparison to the classical Kingman's Coalescent.

1.1 Setup

Let (S, d) be an arbitrary metric space and write $P(S)$ for the space of Borel probability measures on S. We write $P_{\text{fs}}(S)$ for the subspace of finitely supported measures and $P_{\text{cs}}(S)$ for the subspace of compactly supported measures.

We write $C(S)$ for the space of continuous functions on S, $C_b(S) \subset C(S)$ for the subset of bounded continuous functions and $C_0(S) \subset C_b(S)$ for the space of compactly supported functions. For $f \in C_b(S)$ and $\mu \in P(S)$, we write $\mu(f)$ for the integral

$$\mu(f) = \int_S f \, d\mu.$$

We will view $P(S)$ as endowed with the weak topology, which we recall is given by weak convergence of measures; i.e. $\mu_i \to \mu$ weakly if and only if

$$\mu_i(f) \to \mu(f)$$

for all $f \in C_b(S)$. The topology of weak convergence on $P(S)$ is metrizable with the Prokhorov metric d_P. By Prokhorov's theorem, if (S, d) is separable and complete so is $(P(S), d_P)$. [7] [16]

1.2 The Metric Coalescent

Here we will define the Metric Coalescent process $\mu_t, t \geq 0$ from any initial measure μ in $P_{\text{fs}}(S)$. This process is essentially just a reformulation of the Compulsive Gambler process, a toy example of a Finite Markov Information Exchange process (FMIE).[3]

Let $\phi\colon \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ be a fixed continuous function. For a finite measure $\mu \in P_{\text{fs}}(S)$ we write $\#\operatorname{supp}\mu$ for the cardinality of the support of μ. For any such μ, choose arbitrarily a representation of μ as

$$\mu = \sum_{n=1}^{\#\operatorname{supp}\mu} p_i \delta(s_i)$$

where $s_i \in S$ for $1 \leq i \leq \#\operatorname{supp}\mu$ and $p_i = \mu(s_i)$. For $1 \leq i, j \leq \#\operatorname{supp}\mu$ and $i \neq j$ write

$$\nu_{ij} = \phi(d(s_i, s_j)),$$

which will be the instantaneous rate of coalescence between the atoms at s_i and s_j.

The Metric Coalescent process on $P_{\text{fs}}(S)$ is defined as the continuous-time Markov chain with transition rates

$$\mu \to \mu + p_i(\delta(s_j) - \delta(s_i))$$

at rate $\frac{p_i}{p_i+p_j}\nu_{ij}$ for any $1 \leq i, j \leq \#\operatorname{supp}\mu$. Clearly this does not depend on the choice of representation of μ. Equivalently, at rate ν_{ij}, atoms s_i and s_j 'meet' and then merge their mass either all into s_i or s_i with probabilities proportional to p_i and p_j respectively.

Our FMIE interpretation of the MC, by way of the Compulsive Gambler process, is that the sites $s_i \in \operatorname{supp}\mu$ are agents with wealth p_i. These agents then meet at rates according to their geometry in S and upon meeting play a fair winner-take-all game.

The name Metric Coalescent is justified by the following theorem, which follows easily from a comparison with Kingman's Coalescent (as in Section 4) whenever the rate function is non-zero.

Theorem 1.1. *If $\phi > 0$, for any initial measure $\mu_0 \in P_{\text{fs}}(S)$ there is a stopping time $T < \infty$ almost surely, such that μ_T is a point mass distributed as*

$$\mu_T = \delta(\xi)$$

where ξ is distributed as μ_0.

1.3 Extension to $P(S)$

The original Metric Coalescent (MC) process as in Section 1.2 is defined for any metric space (S, d) starting from any finitely supported measure $\mu \in P_{\text{fs}}(S)$. Our goal in this paper is to show for a complete, separable and locally compact metric space (S, d) that this process can be extended to all of $P(S)$. We assume in the sequel unless otherwise stated that our metric space (S, d) satisfies these conditions.

We will also need to make the following two extra assumptions on the continuous rate function ϕ.

(H1) $\phi(x) > 0$ for all $x > 0$.

(H2) $\lim_{x \downarrow 0} \phi(0) = \infty$.

The first assumption will ensure that the process "coalesces". Importantly we do not assume that ϕ is bounded away from zero. A generalization of the MC process without assuming non-zero meeting rates would be of interest; in particular in the Compulsive Gambler process on graphs, pairs of vertices (i.e. agents) only have positive meeting rates if they are connected by an edge [3]. However we do not consider such a generalization in this paper.

The second condition is necessary for the MC to be Feller continuous (see Section 9). Heuristically we think of our rate function as something like $\phi = \frac{1}{x}$

or $\phi = \frac{1}{x^2}$ although the class of allowable rate functions is clearly much more general.

Our main result is the following.

Theorem 1.2. *There exists a unique cadlag Feller continuous $P(S)$-valued Markov process $\mu_t, t \geq 0$ defined from any initial probability measure $\mu_0 = \mu$ on (S, d) such that for any $\mu \in P_{\text{cs}}(S)$*

1. *$\mu_t \in P_{\text{fs}}(S)$ for all $t > 0$, almost surely,*
2. *For each $t_0 > 0$, the process $(\mu_t, t_0 \leq t < \infty)$ is distributed as the Metric Coalescent started at μ_{t_0}.*

There are several variants of Item 2. If the initial measure μ_0 is not compactly supported, then μ_t is still **locally** finitely supported following from 4.1 and the local compactness of S. If $\phi(\cdot)$ is bounded away from zero on $\operatorname{supp} \mu_0$ for some $\mu_0 \in P(S)$ then Theorem 1.1 holds for μ_0 by the same comparison to Kingman's Coalescent. More generally, while any initial measure μ_0 may not every coalesce to a point mass in finite time (for a counterexample see Section 9.2), a simple corollary of the proof of Item 2 is that

$$\lim_{t \to \infty} \mu_t = \delta(\xi),$$

almost surely; that is, the random measure μ_t converges weakly to the random measure given by a point mass chosen from μ_0.

Our method in this paper will be to first construct one such extension of the MC process to all of $P(S)$. We then prove that our construction is Feller continuous which leads to uniqueness in distribution.

We will often refer to our extension as **the** Metric Coalescent, which will be of course justified by Item 2. When we need to distinguish the two, we will often refer to the **original** Metric Coalescent (as defined in Section 1.2) as such.

1.4 Overview

In Section 2, we define the Token Process and its empirical measures and use an exchangeability argument (2.1) to construct the process $\mu_t, t \geq 0$ that will ultimately be our extension of the Metric Coalescent. In Section 3, we prove that the process $\mu_t, t \geq 0$ is a time homogeneous Markov process.

In Section 4, we show that for any compact set $K \subset S$, for all positive times $t > 0$, μ_t has finite support on K. An immediate corollary of this is that for any compactly supported initial measure $\mu \in P_{\text{cs}}(S)$, for all positive times $\mu_t, t > 0$ is finitely supported; in Section 5 it's shown that this process in $P_{\text{fs}}(S)$ is the originally defined Metric Coalescent.

In Section 6 we examine the real-valued processes $\mu_t(f), t \geq 0$ for $f \in C_b(S)$, show that for any f the process is a martingale and use a second moment argument to prove that the MC process is right continuous at $t = 0$ almost surely. In Section 7 we use a coupling argument to prove that MC process is Feller continuous.

Finally, in Section 8 we prove that the MC process is almost surely cadlag and use this to prove that our construction is the unique such extension of the originally defined MC.

We conclude in Section 9 and Section 10 with some motivating examples of the MC process and some open problems and further directions for research.

2 The Token Process

Our construction of the extension of the Metric Coalescent begins with a Markov process we will refer to as the **token process**. For the Compulsive Gambler, an alternative description is provided by augmenting the process with rankings of the agents, where the random winners at meetings between agents are then deterministically replaced by the rankings and their initial randomness. The equivalence in distribution of these two descriptions of the process is a fundamental tool in its study. [3]

Analogously, the token process gives an alternative description of the MC that replaces the mass (or in our heuristic description "wealth") held by agents throughout S with partitions of the natural numbers. The randomness at meeting times is replaced by a simple deterministic model driven by some initial randomness of the process. Most importantly, unlike our original definition of the Metric Coalescent, this alternative description extends naturally to any measure in $P(S)$.

In this section, we first construct the token process and then prove in 2.1 the existence of the limit process $\mu_t, t \geq 0$ that will turn out to be the Metric Coalescent.

2.1 Construction

Fix any initial measure $\mu \in P(S)$. In this section we will define the token process

$$\{(u_i(t), \xi_i)\}_{t\geq 0}^{1\leq i<\infty}$$

generated from μ. For $i \geq 1$ choose ξ_i as IID samples of μ. At time $t = 0$ set $u_i(0) = i$ for all i such that ξ_i is unique among the other samples (the case of non-uniqueness to follow).

For each $1 \leq i, j$ with $i \neq j$ write

$$\nu_{ij} = \phi(d(\xi_i, \xi_j)).$$

We view this as the meeting rate between i and j and write t_{ij} – distributed as an independent rate ν_{ij} exponential random variable – for the meeting time of i, j. If $\xi_i = \xi_j$ – i.e. heuristically the rate ν_{ij} is infinite – then we set $t_{ij} = 0$. The token process will be completely determined by the initial locations $\xi_i, i \geq 1$ and the meeting times t_{ij}, $i \neq j \geq 1$.

From the meeting times $t_{ij}, i, j \geq 1$ we construct a coalescing partition process

$$Z(t) = \{Z_i(t)\}_{i\geq 1, t\geq 0}$$

over the natural numbers $\mathbb{N}$. To distinguish between the elements of this partition from $\mathbb{N}$ used as an index set, we refer to the elements of the copy of $\mathbb{N}$ being partitioned as **tokens**. At time $t = 0$, set

$$Z_i(0) = \{i\}.$$

At each meeting time t_{ij} between $i < j$, if $Z_i \neq \emptyset$ we set

$$Z_i(t) = Z_i(t-) \cup Z_j(t-), \qquad Z_j(t) = \emptyset.$$

If however $Z_i = \emptyset$ or $Z_j = \emptyset$ then we ignore the meeting and call it **trivial**. At any **nontrivial** meeting time t_{ij}, we say that j **merges** with i. Note that each $j \geq 1$ may merge at most once. If a token i has not merged by time t we say that i is **alive** at time t.

For $N \geq 1$ we will also consider the N-th partition processes

$$Z^N(t) = \left(Z_1^N(t), \ldots, Z_N^N(t)\right)$$

given by

$$Z_n^N(t) = Z_n(t) \cap [\mathrm{N}],$$

where $[\mathrm{N}] = \{1, 2, \ldots, N\}$.

In the case of two or more tokens initially occupying the same site – i.e. whenever $\xi_i = \xi_j$ for some $i \neq j$ – we merge all the tokens at that site instantaneously into the 'lowest' token. This is consistent if more than two tokens occupy the same site, since if

$$t_{ij} = t_{ij'} = 0$$

then also $t_{jj'} = 0$. The end result is that for any site $s \in S$, the lowest token $\bar{i}$ at s starts at $t = 0$ owning all the other tokens at s, i.e.

$$Z_{\bar{i}}(0) = \{j \colon \xi_j = \xi_{\bar{i}}\}$$

and $Z_j(0) =$ for all other j with $\xi_j = \xi_{\bar{i}}$.

Write $u_i(t)$ for the owner of token i at time t, that is define $u_i(t)$ by

$$i \in Z_{u_i(t)}.$$

The path of token i transitions as follows: if $u_i(t_{jk}-) = j$ and $k < j$ is still alive at time t_{jk}, then at time t_{jk} we have

$$u_i(t_{jk}) = k.$$

Thus, the path of each token i – i.e. $u_i(t), t \geq 0$ – depends only on the meeting times amongst tokens 1 to i, showing that the infinite token process can be defined consistently.

2.2 The Empirical Measures

Our main objects of focus to begin will be the random empirical measures μ_t^N, given by

$$\mu_t^N = \frac{1}{N} \sum_{i=1}^{N} \delta(\xi_{u_i(t)}),$$

or equivalently

$$\mu_t^N = \sum_{i=1}^{N} \frac{\#Z_i^N(t)}{N} \delta(\xi_i).$$

From the N-token process, via a limiting argument, we will define the stochastic process $\mu_t, t \geq 0$ in $P(S)$ that will be shown to be the extension of the Metric Coalescent.

Proposition 2.1. *For any initial measure $\mu \in P(S)$, there exists a random measure-valued process $\mu_t, t \geq 0$ with $\mu_0 = \mu$ such that for all $t \geq 0$*

$$\mu_t^N \to \mu_t, \text{ almost surely.}$$

In this section our main goal will be to prove this proposition.

2.3 Exchangeable Partitions, Measure-Valued Markov Processes

Throughout we'll rely heavily on the theory of exchangeable partitions, which we will set up. This account closely follows that in [6].

A partition π of $\mathbb{N}$ can be thought of either through its blocks or by the equivalence relation $\sim^\pi$ it induces on $\mathbb{N}$; i.e. n and m are in the same block if and only if

$$n \sim^\pi m.$$

For any permutation σ on $\mathbb{N}$, we can consider the partition $\sigma\pi$ given by

$$i \sim^{\sigma\pi} j$$

if and only if

$$\sigma(i) \sim^\pi \sigma(j).$$

Writing Π for the space of partitions of $\mathbb{N}$, a random partition $\pi \in \Pi$ is called **exchangeable** if for any finite permutation σ of $\mathbb{N}$ - i.e. constant except on finitely many $n \in \mathbb{N}$ - π and $\sigma\pi$ have the same distribution on Π.

As defined in Section 2.1, the token process gives rise to the random family of partitions $Z(t), t \geq 0$ on $\mathbb{N}$, which we can specify by their equivalence relation $\sim_t, t \geq 0$. Importantly, note that for two tokens i and j, we have

$$i \sim_t j$$

if and only if the locations of i and j are the same at time t; i.e.

$$\xi_{u_i(t)} = \xi_{u_j(t)}.$$

To the astute reader familiar with measure-valued Markov processes, our description of the token process is superficially very similar to the Donnelly-Kurtz Lookdown process. [2] The difference of course is in the details; in our setting mergers do not have corresponding birth events. More importantly, the mechanism by which tokens (corresponding loosely to levels) are selected is both dependent on the empirical measures and chosen by an exchangeable, geometrically dependent (i.e. thus not independent) meeting clock.

A closer analogy can be found between the token process and Kingman's Coalescent. In fact, this analogy will be made explicit (see Section 4.1) later on, so we'll suffice for now to note that the token process can in fact be thought of as a generalization of Kingman's model, however with (geometrically defined) exchangeable rates in place of independent rates.

2.4 Two Types of Exchangeability

The proof of 2.1 will rely upon two different forms of exchangeability for the token process, corresponding to two different descriptions of the limit process $\mu_t, t \geq 0$. First, we have what we will call the **symmetric exchangeability** which will show that at any time t, conditional on the past (of the measure process), the token partition $Z(t)$ at time t is an exchangeable partition.

The second type of exchangeability needed in our proof of 2.1 we'll call **asymmetric**. This will show that for any time $t \geq 0$, conditional on the location of the first K tokens, the location of the subsequent tokens are also exchangeable.

Both of these ultimately imply that at time t, the partition $Z(t)$ is exchangeable. However, both also say a bit more about that exchangeability, and combined together lead to our proof of 2.1.

2.4.1 Symmetric Exchangeability

The first form of exchangeability is a bit harder to fully describe without making reference to $\mu_t, t \geq 0$ which currently still needs to be constructed. A more complete view of this exchangeability is given in 3.3.

Our main goal here is to demonstrate an approximation of this exchangeability (2.3) and use this to prove the following.

Proposition 2.2. *From any initial measure μ and for all times $t \geq 0$, the random partition process $Z(t)$ is exchangeable almost surely.*

Recall that to show this, we'll need to show that for any time $t \geq 0$ and any finite permutation σ, the distributions of π_t and $\sigma\pi_t$ are the same. To prove 2.2, we'll focus on the changing partitions Z^N of $[\mathrm{N}] = \{1, 2, \ldots, N\}$ and prove that this gives an exchangeable partition of [N] for $N \geq 1$, which is clearly sufficient.

Interestingly, this proof - and thus the exchangeability of the whole partition process - relies only on how the blocks of the partition merge and not in any way on the particulars of the meeting structure (only that these meetings are exchangeable). This gives a nice example of the motivation behind the dichotomy in the framework of FMIE processes, separating the structure of a particle system into the meeting model and the information exchange model.

Let $r_1, \ldots, r_N$ be a uniform random ordering of $1, \ldots, N$. Recall the partition process $Z_i^N(t)$ defined by the token process and write $S_i^N(t)$ for

$$S_i^N(t) = Z_{r_i}^N(t),$$

for $t \geq 0$, $1 \leq i \leq N$; i.e. for the tokens owned by r_i at time t. Write $S^N(t)$ for the random partition process

$$S^N(t) = (S_1^N(t), \ldots, S_N^N(t)),$$

of [N].

For any partition α of [N], write $|\alpha|$ for the trace $(\#\alpha_1, \ldots, \#\alpha_K)$ of the partition. We will often not specify the size, i.e. number of blocks, of a partition as in the context it will be clear.

Note that if for some i, j we have $\xi_{r_i} = \xi_{r_j}$, then by construction $t_{r_i r_j} = 0$ and so the merger of blocks S_i^N and S_j^N has "already" occurred at time $t = 0$.

Thus, as opposed to the standard set-up of coalescing partition processes, we don't necessarily begin from a partition of all singletons. Unfortunately this serves to confuse the notation a bit but is ultimately irrelevant to our proof.

Proposition 2.3. *For any $t_0 > 0$, the conditional distribution of $S^N(t_0)$ given $\mu_t^N, 0 \leq t \leq t_0$ is uniform over all possible partitions. That is, for all α, β with $|\alpha| = |\beta|$ we have*

$$\mathbb{P}(S^N(t_0) = \alpha | \mu_t^N, 0 \leq t \leq t_0) = \mathbb{P}(S^N(t_0) = \beta | \mu_t^N, 0 \leq t \leq t_0).$$

Proof. Instead of looking at the continuous time partition process, we will consider instead the discrete jump process. Write $T_0 = 0$ and let $T_n, n \geq 1$ be the time of the n-th (non-trivial) meeting, i.e. the n-th jump of the process. Write $\mathcal{F}_n = \sigma(|S^N|(n))$ for $n \geq 0$. We will show that for α, β with $|\alpha| = |\beta|$

$$\mathbb{P}(S^N(n) = \alpha | \mathcal{F}_n, \ldots, \mathcal{F}_0) = \mathbb{P}(S^N(n) = \beta | \mathcal{F}_n, \ldots, \mathcal{F}_0)$$

for any $n \geq 0$, which is clearly equivalent to the continuous time claim. Note that F_0 is not trivial, as possibly singleton blocks are already merged at time $t = 0$.

Conditioning on $\mathcal{F}_n, \ldots, \mathcal{F}_0$, we know at each (non-trivial, positive) meeting which two blocks of the partition S^N have merged and which of those blocks has the lowest ranking token - i.e. the "winner" of the meeting - but not what these tokens are. For $1 \leq k \leq n$, define i_k and j_k by the rule that the partition $S^N(k)$ follows from $S^N(k-1)$ by merging the i_k block into the j_k block; that is $S_{i_k}(k) = \emptyset$ and

$$S_{j_k}^N(k) = S_{j_k}^N(k-1) \cup S_{i_k}^N(k-1).$$

Equivalently at time k we have i_k meeting and losing to j_k.

From any initial configuration $S^N(0) = \gamma$, define $f_\gamma \colon [\mathrm{N}] \to [\mathrm{N}]$ by

$$f_\gamma(k) = \inf S_k^N(0)$$

for the owner of block k, noting the convention that if $S_k^N = \emptyset$ then $f_\gamma(k) = \infty$.

Call an initial configuration γ **compatible** if for all meetings $1 \leq k \leq n$ we have

$$f_\gamma(j_k) < f_\gamma(i_k),$$

and say the initial configuration γ **results** in α if

$$\{S^N(0) = \gamma\} \cap \bigcap_{k=1}^{n} \{k - \text{th meeting is between } i_k \text{ and } j_k\} \subset \{S^N(n) = \alpha\}.$$

By our setup, $S^N(0)$ is uniform over all possible initial configurations - i.e. initial partitions γ with $|\gamma| = |S^N(0)|$. Thus $\mathbb{P}(S^N(n) = \alpha | F_n, \ldots, F_0)$ is the fraction of compatible initial configurations that result in $S^N(n) = \alpha$.

Instead of counting this directly, consider a fixed time n partition α. For any other partition β with $|\beta| = |\alpha|$ we will simply show a bijection between compatible initial configurations resulting in each. Fix such a β.

Let τ be the unique bijection on $\{1, \ldots, N\}$ which for each $1 \leq i \leq N$ maps α_i to β_i in an order preserving fashion; i.e. for each $1 \leq i \leq N$ the first element

of α_i maps to the first of β_i, the second element of α_i maps to the second of β_i etc. As we know $\#\alpha_i = \#\beta_i$ for each i this is well defined.

If γ results in α, we claim the initial configuration $\tilde{\gamma}$ given by

$$\tilde{\gamma}_i = \tau(\gamma_i)$$

results in β. To check that $\tilde{\gamma}$ is compatible: after each meeting of blocks i_k and j_k, $f_\gamma(i_k)$ and $f_\gamma(j_k)$ end up in the same block of the partition, and therefore at time n are still in the same block α_i (for some i). Therefore as $f_\gamma(j_k) < f_\gamma(i_k)$ we also have $f_{\tilde{\gamma}}(j_k) < f_{\tilde{\gamma}}(i_k)$ since τ is order preserving within each block of α.

Therefore, we find that for arbitrary α, β with $|\alpha| = |\beta| = |S^N|(n)$ that

$$\mathbb{P}(S^N(n) = \alpha | F_n, \dots, F_0) = \mathbb{P}(S^N(n) = \beta | F_n, \dots, F_0).$$

If $|\alpha| = |\beta| \neq |S^N|(n)$ then the equality is trivially true as both are zero. □

From 2.3, noting that the probability of any given permutation doesn't depend on the labelling r_k of the blocks, 2.2 immediately follows.

2.4.2 Asymmetric Exchangeability

The second, asymmetric exchangeability property of the token process is exchangeability in the sense of de Finnetti's theorem. For it, we consider a fixed token K and time $t \geq 0$ and look at the locations of the tokens $i \geq K+1$. Recall that $u_i(t)$ is the owner of token i at time t. The asymmetric exchangeability is given by the following.

Proposition 2.4. *For any fixed $K \geq 1$, $t \geq 0$, the sequence of indicators*

$$1(u_{K+1}(t) = K), 1(u_{K+2}(t) = K), 1(u_{K+3}(t) = K), \dots$$

is, conditional on $\xi_1, \dots, \xi_K$, an infinitely exchangeable sequence.

Note that both 2.4 and 2.3 ultimately imply 2.2, i.e that the token partition is exchangeable. The key difference is in where this exchangeability is coming from and what it is conditional on. The symmetric exchangeability is in relation to the history of the limiting process μ_t, in relation to which (as we will see in 3.3) the different tokens behave symmetrically. The exchangeability here in 2.4 is clearly asymmetric with respect to the tokens. The relationship between these two viewpoints is key in the sequel.

Our proof of 2.4 will make use of the combinatorial tool of **meeting trees** of the token process. For a fixed $N \geq 1$, we define a meeting tree T of the first N tokens as an ordered list

$$T = (a_1, a_2, \dots; b_1, b_2, \dots | i_1 \to j_1, \dots, i_n \to j_n)$$

which recounts the order of the meetings of the first N tokens but not the (non-zero) times of these meetings.

We read the tree as follows. To the left of the divide are groups of tokens starting at the same initial location and thus meeting at time $t = 0$; i.e. for the generic tree T above

$$\xi_{a_1} = \xi_{a_2} = \dots, \xi_{b_1} = \xi_{b_2} = \dots,$$

and so each such group meets at time $t = 0$ and merges into the lowest token of the group. We make no attempt at "ordering" meetings that occur simultaneously at time zero.

On the right of the divide, we write $a \to b$ for the meeting of tokens a and b resulting in a merging into token b. We call a tree **valid** if

1. For each meeting $i_r \to j_r$, we have $1 \leq i_r < j_r \leq n$.
2. Each token 1 to N can merge into another token at most once, including meetings at $t = 0$.

Of course each token can be *merged into* multiple times in the history of the process.

For any $t \geq 0$, write $S^N(t)$ for the (random) observed meeting tree at time t of the N token process. Then for any valid meeting tree T, the probability of $S^N(t) = T$ can be explicitly calculated as

$$\mathbb{P}(S^N(t) = T | \xi_1, \ldots, \xi_N) = f_T(\xi_1, \ldots, \xi_N)$$

for some function $f_T \colon S^N \to \mathbb{R}$ of the initial locations of the tokens. We'll see later (see Section 6.2 and Section 7.5) that for all but the smallest N calculating f_T explicitly is impractical.

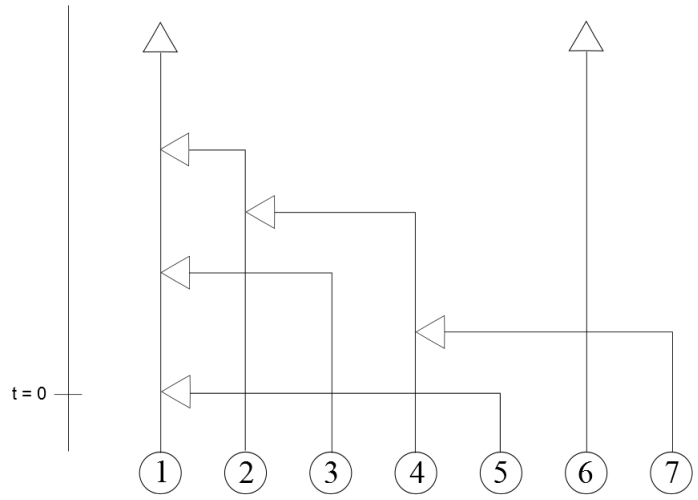

Figure 1: Meeting Tree $T = (5, 1 | 7 \to 4, 3 \to 1, 4 \to 2, 2 \to 1)$

We are now ready to give a proof of 2.4.

Proof. Fix $K \geq 1$ and $t \geq 0$. To show that the sequence

$$1(u_{K+1}(t) = K), 1(u_{K+2}(t) = K), 1(u_{K+3}(t) = K), \ldots$$

is conditionally exchangeable, it suffices to show that for any $N > M \geq K + 1$ that the conditional distribution of

$$V = \Big(1(u_{K+1}(t) = K), \ldots 1(u_M(t) = K), 1(u_{M+1}(t) = K), \ldots 1(u_N(t) = K)\Big)$$

is the same as the conditional distribution of

$$\hat{V} = \Big(1(u_{K+1}(t) = K), \dots 1(u_{M+1}(t) = K), 1(u_M(t) = K), \dots 1(u_N(t) = K)\Big)$$

i.e. under a transposition of M and $M+1$, since the set of transpositions generates the group of all finite permutations.

Consider any fixed binary sequences s' and s'' of length $M-K$ and $N-(M-K+2)$ respectively. Then for any binary sequence s of length $N-K$ of the form

$$s = (s', *, **, s'')$$

where $*, ** \in \{0, 1\}$ it suffices to show that

$$\mathbb{P}\left(V = (s', *, **, s'')|\xi_i, 1 \le i \le K\right) = \mathbb{P}\left(\hat{V} = (s', *, **, s'')|\xi_i, 1 \le i \le K\right)$$

or equivalently

$$\mathbb{P}\left(V = (s', *, **, s'')|\xi_i, 1 \le i \le K\right) = \mathbb{P}\left(V = (s', **, *, s'')|\xi_i, 1 \le i \le K\right).$$

Now, if $*$ and $**$ are both either 0 or 1 then this is trivially true, so we need only consider the case $* \neq **$. Without loss of generality assume $* = 0$ and $** = 1$. Let τ be the set of all meeting trees resulting in $V = (s', 0, 1, s'')$ and τ' the set of all meeting trees resulting in $V = (s', 1, 0, s'')$. Note that both τ and τ' are finite as there are only finitely many possible meeting trees on N tokens.

We will construct a bijection $\psi\colon \tau \to \tau'$ with the property that for all $T \in \tau$

$$\mathbb{P}(S^N(t) = T|\xi_i, 1 \le i \le K) = \mathbb{P}(S^N(t) = \psi(T)|\xi_i, 1 \le i \le K),$$

which will show that

$$\begin{aligned}
&\mathbb{P}\Big(V = (s', 0, 1, s'')|\xi_i, 1 \le i \le K\Big) \\
&\qquad = \sum_{T \in \tau} \mathbb{P}\left(S^N(t) = T|\xi_i, 1 \le i \le K\right) \\
&\qquad = \sum_{T \in \tau} \mathbb{P}\left(S^N(t) = \psi(T)|\xi_i, 1 \le i \le K\right) \\
&\qquad = \sum_{T \in \tau'} \mathbb{P}\left(S^N(t) = T|\xi_i, 1 \le i \le K\right) \\
&\qquad = \mathbb{P}\left(V = (s', 1, 0, s'')|\xi_i, 1 \le i \le K\right).
\end{aligned}$$

For any $T \in \tau$, we define $\psi(T)$ to be the meeting tree given by switching the paths of the tokens M and $M+1$. A priori this need not be a valid meeting tree. In particular, if $M+1 \to M$ occurs in T, then $\psi(T)$ would have the meeting $M \to M+1$ and so wouldn't be a valid meeting tree. However, by assumption $T \in \tau$ has

$$1(u_M(t) = K) = 0, \qquad 1(u_{M+1}(t) = K) = 1$$

and so at time t token $M+1$ is owned by K and token M is not. Thus the paths of the tokens M and $M+1$ can not have already met in T, as otherwise either both would be owned by K or neither would be. Therefore $\psi(T)$ is a valid meeting tree and so is clearly in τ'. To see that ψ is a bijection, we need

only note that reversing the paths of M and $M+1$ again gives the inverse $\psi^{-1}\colon \tau' \to \tau$.

To see T and $\psi(T)$ have the same conditional probability we need only note that ξ_M and ξ_{M+1} are exchangeable and

$$f_{\psi(T)}(\xi_1, \ldots, \xi_N) = f_T(\xi_1, \ldots, \xi_{M+1}, \xi_M, \ldots \xi_N)$$

and so

$$\begin{aligned}\mathbb{E}\left(f_{\psi(T)}(\xi_1, \ldots, \xi_N)|\xi_i, 1 \leq i \leq K\right) \\ &= \mathbb{E}\left(f_T(\xi_1, \ldots, \xi_{M+1}, \xi_M, \ldots, \xi_N)|\xi_i, 1 \leq i \leq K\right) \\ &= \mathbb{E}\left(f_T(\xi_1, \ldots, \xi_N)|\xi_i, 1 \leq i \leq K\right)\end{aligned}$$

completing the proof. □

In fact our proof shows quite a bit more. Fixing a $K \geq 1$ and a time $t \geq 0$, for $n \geq K+1$ consider the vector V_n given by

$$V_n = \left(1(u_n(t) = 1), 1(u_n(t) = 2), \ldots, 1(u_n(t) = K)\right).$$

It follows easily from our proof of 2.4 that conditional on the entire history of the first K tokens up to time t - i.e. on $\xi_1, \ldots, \xi_K$ and $Z^K(s), 0 \leq s \leq t$ that

$$V_{K+1}, V_{K+2}, V_{K+3}, \ldots$$

is an infinite exchangeable sequence. The value of 2.4 in the sequel is the ability to do computations by conditioning on only the initial conditions of finitely many tokens and so we content ourself with the current statement of the result.

2.5 Proof of 2.1

We are now ready to start our proof that the empirical measures μ_t^N coming from the token process converge weakly for each $t \geq 0$. For time $t = 0$, this is essentially the Glivenko-Cantelli theorem.

Lemma 2.5. *$\mu_0^N \to \mu$ almost surely and thus $\mu_0 = \mu$ a.s.*

We note here however that while the convergence is easy, actually quantifying this convergence (in the Wasserstein metric, for instance) is hard and to our knowledge impossible in the absence of finite-dimensional type assumptions on (S, d) [8]. More classically for measures on $\mathbb{R}^d$ a bit more is known [14]. Interestingly, more can be said for the MC process at positive times $t > 0$, and we discuss later qualitative estimates of this convergence for positive times $t > 0$ in Section 2.6.

We focus on showing weak convergence for times $t > 0$. Our main tool in this in Kingman's Paintbox Theorem, which serves as a structure theorem for exchangeable partitions.

2.5.1 Kingman's Paintbox Theorem

The key application of 2.2 is Kingman's Paintbox Theorem. Following Theorem 1.1 in [6], a paintbox partition is derived from a mass partition $p = (p_0, p_1, \ldots)$ with $p_i \geq 0$ satisfying

$$\sum_{i=0}^{\infty} p_i = 1,$$

and

$$p_1 \geq p_2 \geq \ldots.$$

Note that we do not require $p_0 \geq p_1$. Write $\mathfrak{P}_{\mathrm{m}}$ for the set of all such mass partitions.

Let $U_1, U_2, \ldots$ be an I.I.D. sequence of uniform $(0,1)$ random variables and write

$$I(u) = \inf\{n \geq 0 \colon \sum_{m=0}^{n} p_i > u\}.$$

Intuitively, if p is a partition of the unit interval, then $I(u)$ is the block of partition that u falls into. Then, the paintbox partition given by p is defined by $i \sim j$ if and only if

$$I(U_i) = I(U_j) \geq 1,$$

i.e. if U_i and U_j fall in the same (non p_0) "slot" of the paintbox. The mass p_0 corresponds to singletons blocks of the partition, called **dust**. An elementary fact of exchangeable partitions is that all blocks must be either singletons or infinite, and don't take on any other finite sizes.

Clearly this construction gives an exchangeable partition. The importance of Kingman's Theorem is that this is a universal construction for such random partitions. Our statement of the theorem follows Theorem 1.1 in [6].

Theorem 2.6. *Kingman's Paintbox: Let π be any exchangeable partition. Then there exists a probability measure $\alpha(dp)$ on $\mathfrak{P}_{\mathrm{m}}$ such that*

$$\mathbb{P}(\pi \in \cdot) = \int_{p \in \mathfrak{P}_{\mathrm{m}}} \alpha(dp) \rho_p(\cdot),$$

where for a mass partition p, ρ_p is the corresponding law on Π given by the paintbox construction.

2.5.2 Defining μ_t

Applying Kingman's correspondence, the token process' partition $Z(t)$ can be described as follows. First, there exists a (random) mass partition $p(t) = \{p_k(t), k \geq 0\}$ with

$$\sum_{k=0}^{\infty} p_k(t) = 1, \tag{1}$$

giving a random partition

$$B(t) = \{B_k(t), k \geq 1\},$$

of $\mathbb{N}$, which will have asymptotic rates $p_k(t)$. Note that the partition $B(t)$ will be the same as $Z(t)$, however is no longer indexed by the lowest token in each block.

Then, the equivalence relation generating the token partition $\sim_t$ is given by

$$i \sim_t j$$

if and only if i and j are in the same block $B_k(t)$ for $k \geq 1$.

Conditional on $p(t)$, each token $i \geq 1$ is assigned a block of $B(t)$ independently according to $p(t)$. That is, token i is assigned to block B_i with probability $p_i(t)$ for $i \geq 1$. The remaining mass $p_0(t)$ corresponds to token i being a singleton (or dust), which in our setting would be tokens that have had no meetings by time t. There are many different ways to show that this doesn't happen almost surely; our proof uses the asymmetric exchangeability of 2.4.

Lemma 2.7. *For any $t > 0$, almost surely $p_0(t) = 0$.*

Proof. By construction, $\mathbb{E}\, p_0(t)$ is the probability that any given token is a singleton at time t, in particular we'll consider the probability that token 1 is a singleton at time t. The key here is to calculate this using the second form of exchangeability.

By 2.4 the sequence

$$1(u_2(t) = 1), 1(u_3(t) = 1), \ldots$$

is conditionally exchangeable given ξ_1. Applying de Finetti's theorem to this sequence of Bernoulli variables (Theorem 4.6.6, [10]) we see that there is a $q_1(t) \in \sigma(\xi_1)$ satisfying $0 \leq q_1(t) \leq 1$ such that conditional on $q_1(t)$, the sequence is independent and identically distributed (as $\mathrm{Ber}(q_1(t))$ random variables).

On the other hand, by the definition of the token process we have that

$$\begin{aligned}
\mathbb{P}(u_2(t) = 1|\xi_1) &= \mathbb{P}(t_{12} \leq t|\xi_1) \\
&= \mathbb{E}\left(1 - \exp(-\phi(d(\xi_1, \xi_2))t)|\xi_1\right) \\
&> 0,
\end{aligned}$$

almost surely for all $t > 0$. Therefore, we have that

$$q_1(t) > 0$$

almost surely.

Since $1(u_i(t) = 1)$ are independent conditional on $q_1(t)$, by the Strong Law they can only all be zero if $q_1(t)$ is itself zero. Thus, we can conclude that

$$\begin{aligned}
\mathbb{P}(1 \text{ is dust at } t) &= \mathbb{P}\left(\forall i \geq 2, 1(u_i(t) = 1) = 0\right) \\
&= \mathbb{E}\,\mathbb{P}\left(\forall i \geq 2, 1(u_i(t) = 1) = 0|\xi_1\right) \\
&= \mathbb{E}\,\mathbb{P}(q_1(t) = 0|\xi_1) \\
&= 0
\end{aligned}$$

whenever $t > 0$, completing the proof. □

We are now finally able to define the random measure μ_t. By construction of the token process, two tokens i and J have $i \sim_t j$ if and only if they are at the same location of S at time t. By Lemma 2.7, this gives us a one-to-one correspondence between blocks $B_k(t)$ and locations $s_k \in S$. Namely, writing

$$l_k(t) = \inf B_k(t)$$

for the lowest token in block B_k, we have

$$s_k(t) = \xi_{l_k(t)}.$$

Therefore we can define our random measure μ_t by

$$\mu_t = \sum_{k=1}^{\infty} p_k(t)\delta(s_k(t)). \tag{2}$$

By Lemma 2.7 and Equation (1) we are guaranteed that μ_t is in fact a probability measure on S.

All that remains of 2.1 is to show that for all $t > 0$

$$\mu_t^N \to \mu_t$$

almost surely.

2.5.3 Convergence

Later (Section 2.6) we give of a proof of the convergence of the empirical measures using Total Variation distance. However, we'll need to make use later of the concept of a convergence determining class for weak convergence and so we opt here to give such a proof of 2.1.

For any $k \geq 1$, write

$$p_k^N(t) = \frac{1}{N}\sum_{n=k}^{N} 1(n \in B_k(t))$$

for the empirical mass at s_k of the first N tokens. We first show that for any function $f \in C_b(S)$, the measures μ_t^N converge to μ_t almost surely.

Lemma 2.8. *For any fixed $f \in C_b(S)$ and for all $t > 0$,*

$$\mu_t^N(f) \to \mu_t(f)$$

almost surely.

Proof. For $i \geq 1$ write $f_i, i \geq 1$ for the sequence

$$f_i(t) = \sum_{i=k}^{\infty} f(s_k(t))1(i \in B_k(t)).$$

By assumption $f_i(t)$ is bounded, since f is and only one such indicator occurs. Clearly

$$\mu_t^N(f) = \frac{1}{N}\sum_{i=1}^{N} f_i(t).$$

Recalling the Paintbox construction, conditional on $p(t)$ the sequence f_i is independent and so applying the Strong Law we can conclude that

$$\mu_t^N(f) \to \mu_t$$

almost surely. □

To complete the proof of 2.1 we need to show that $\mu_t^N(f) \to \mu_t(f)$ for all $f \in C_b(S)$, almost surely. To show this from Lemma 2.8, we use a standard reduction taking advantage of the fact that S is separable.

The key is the standard result that there exists a countable class of test functions for weak convergence, called a **convergence determining class** [7]. This is opposed to needing to test all of $C_b(S)$ for convergence, which is (for generic S) uncountable. Since we'll apply this same reduction a few times in the sequel, we expand on the details here.

Theorem 2.9. *If S is a separable metrizable space, then there exists a countable subset $\mathfrak{C} \subset C_b(S)$ such that for any sequence of measures $\nu^i, 1 \le i \le \infty$, $\nu^i \to \nu^\infty$ weakly if and only if $\nu^i(f) \to \nu^\infty(f)$ for all $f \in \mathfrak{C}$.*

An immediate application to our setting is the following corollary.

Proposition 2.10. *If S is a separable metrizable space, then for any sequence of random measures $\nu^i \in P(S), 1 \le i \le \infty$ defined on the same probability space, $\nu^i \to \nu^\infty$ in the weak topology almost surely if and only if for every fixed $f \in C_b(S)$, $\nu^i(f) \to \nu^\infty(f)$ almost surely.*

Applying 2.10 and Lemma 2.8 the proof of 2.1 us complete.

2.6 Support and Total Variation Distance

We conclude this section with a brief discussion of the support of the measures $\mu_t, t \ge 0$ and give an easy corollary of 2.1.

Note that the measures μ_t^N have finitely many atoms and that their support coincides with the locations at time t of the first N tokens, that is

$$\operatorname{supp} \mu_t^N = \{\xi_{u_i(t)} \colon 1 \le i \le N\}.$$

For μ_t, by an abuse of notation for now we'll define the support for $t > 0$ as the similarly defined set of token's locations, i.e. we'll define $\operatorname{supp} \mu_t$ by

$$\operatorname{supp} \mu_t = \{\xi_{u_i(t)} \colon i \ge 1\}. \tag{3}$$

We'll show later (see Section 4) that for positive times the set of token locations is finite on every compact set of S, and so is nowhere dense. This implies that our definition is in fact equivalent to the standard definition of the support of the measure, i.e

$$\operatorname{supp} \nu = \overline{\{s \in S \colon \nu(U) > 0 \text{ for all open } U \ni s\}},$$

as shown in Lemma 4.3. In particular, this shows that for all positive times t, the support of μ_t is countable. An immediate and useful corollary of this is the following.

Lemma 2.11. *For all $t > 0$*

$$\operatorname{supp} \mu_t^N \subset \operatorname{supp} \mu_t.$$

We can now prove the following strengthening of 2.1 which we will need later, showing that μ_t^N converges to μ_t in the stronger topology of Total Variation distance. We recall the definition of the Total Variation distance for a countable state space[12]: for two measures m_1 and m_2 on a countable state space X, the total variation distance d_{TV} is given by

$$d_{\mathrm{TV}}(m_1, m_2) = \frac{1}{2} \sum_{x \in X} |m_1(x) - m_2(x)|.$$

As the token measures μ_t^N and their limit μ_t share a countable state space, by a fairly standard argument we can calculate their Total Variation distance.

Lemma 2.12. *For any $t > 0$, we have*

$$d_{\mathrm{TV}}(\mu_t^N, \mu_t) \to 0$$

almost surely.

Proof. Fix any $\epsilon > 0$. Then, because

$$\sum_{k=1}^{\infty} p_k(t) = 1$$

almost surely, we can almost surely choose a K so that

$$\sum_{k=K+1}^{\infty} p_k(t) < \epsilon. \tag{4}$$

Next, the Strong Law applied to the paintbox construction gives that $p_k^N(t) \to p_k(t)$ for all $k \geq 1$ almost surely. Thus, we can choose almost surely an M such that if $N \geq M$ then

$$\sum_{k=1}^{K} |p_k^N(t) - p_k(t)| \leq \epsilon.$$

We claim that

$$\sum_{k=K+1}^{\infty} p_k^N(t) < 2\epsilon.$$

To see this, we calculate

$$\begin{aligned}
|\sum_{K+1}^{\infty} p_k^N(t) - \sum_{K+1}^{\infty} p_k(t)| &= |\left(1 - \sum_{1}^{K} p_k^N(t)\right) - \left(1 - \sum_{1}^{K} p_k(t)\right)| \\
&= |\sum_{1}^{K} p_k^N(t) - \sum_{1}^{K} p_k(t)| \\
&\leq \sum_{k=1}^{K} |p_k^N(t) - p_k(t)| \\
&\leq \epsilon,
\end{aligned}$$

which combined with Equation (4) proves the claim.

Thus, for $N \geq M$, we have that

$$\begin{aligned} d_{\mathrm{TV}}(\mu_t^N, \mu_t) &= \frac{1}{2}\sum_{k=1}^{\infty} |p_k^N(t) - p_k(t)| \\ &\leq \frac{1}{2}\sum_{k=1}^{K} |p_k^N(t) - p_k(t)| + \frac{1}{2}\sum_{K+1}^{\infty} \left(p_k^N(t) + p_k(t)\right) \\ &\leq \frac{5}{2}\epsilon, \end{aligned}$$

completing the proof. □

Note that Lemma 2.12 is trivially false for time $t = 0$. In particular, if μ is non-atomic, as all of the empirical measures μ_0^N are atomic, they can never converge to $\mu_0 = \mu$ in Total Variation distance.

Note that in light of Lemma 2.12 it seems natural to wonder if μ_t^N converges to μ_t uniformly in Total Variation distance for $t > 0$. We answer this question in the affirmative for t bounded away from 0 in Lemma 8.2.

3 The Markov Property

In this section, we'll prove that from any initial measure $\mu \in P(S)$, that the constructed process $\mu_t, t \geq 0$ satisfies the Markov Property.

Proposition 3.1. *From any initial μ, the process $\mu_t, t \geq 0$ is a time-homogeneous Markov process.*

Our approach to this will be to show first that the empirical measures $\mu_t^N, t \geq 0$ of the token process are Markovian with respect to the adapted filtration of $\mu_t, t \geq 0$. Then, by a limiting argument we will show that this carries over for μ_t. Finally, we use the exchangeability of the token process to prove time homogeneity.

3.1 Markov Condition for μ_t^N

To begin, we show that the empirical processes $\mu_t^N, t \geq 0$ are Markov with respect to the filtration generated by the MC process $\mu_t, t \geq 0$.

Proposition 3.2. *For any Borel $A \subset P(S)$, times $0 \leq t_0 \leq t_1$ and $N \geq 1$ we have*

$$\mathbb{P}(\mu_{t_1}^N \in A | F_{t_0}) = \mathbb{P}(\mu_{t_1}^N \in A | \mu_{t_0}).$$

To prove 3.2, we fix a time t_0 and make use of the symmetric paintbox description of the measure μ_{t_0} as

$$\mu_{t_0} = \sum_{k=1}^{\infty} p_k(t_0)\delta(s_k).$$

Now each token $1 \leq i \leq N$ at time t_0 is in one of the blocks of the paintbox partition given by $p(t_0)$, so write A_i for the block of the paintbox partition containing token i. That is, the location in S of token i at t_0 is

$$\xi_{u_i(t_0)} = s_{A_i}.$$

For simplicity of notation we will often write

$$A^N = (A_1, \ldots, A_N),$$

for the location of the first N tokens at time t_0. By 2.3, we already know that at time t_0 the token partition $Z(t_0)$ is exchangeable. In fact, we will show that even conditional on $\mu_t, 0 \leq t \leq t_0$, the token process is still exchangeable. This is the complete description of the symmetric form of exchangeability for the MC process (as in Section 2.4.1).

Proposition 3.3. *Conditional on $\mu_t, 0 \leq t \leq t_0$ the token partition $Z(t_0)$ is exchangeable.*

Proof. 2.3 shows that conditional on $\mu_t^N, 0 \leq t \leq t_0$ the partition $Z^N(t_0)$ is exchangeable. To complete this proof, we need to show that this exchangeability holds conditional on $\mu_t, t \geq 0$.

A (standard) basis for the topology of the set Π of all permutations on $\mathbb{N}$ is given by the finite restrictions, that is by the sets of the form

$$B_{\pi_0} = \{\pi \in \Pi \colon \pi|_N = \pi_0|_N\}$$

where $\pi|_N = \pi \cap \{1, 2 \ldots, N\}$. Consider any such fixed $B = B_{\pi_0}$.

Let σ be a finite permutation of $\mathbb{N}$. By assumption there is an $N_\sigma \geq 1$ such that σ is constant for all integers after N_σ. Now, 2.3 implies that the conditional law of the partition $Z^N(t_0)$ on the set Π of all positive partitions satisfies

$$\mathbb{P}(\sigma Z^N(t_0) \in B | \mu_t^N, 0 \leq t \leq t_0) = \mathbb{P}(Z^N(t_0) \in B | \mu_t^N, 0 \leq t \leq t_0),$$

for all $N \geq N_\sigma$.

To show that $Z(t_0)$ is conditionally exchangeable (on B), we need to show that for any such σ,

$$\mathbb{P}(\sigma Z(t_0) \in B | \mu_t, 0 \leq t \leq t_0) = \mathbb{P}(Z(t_0) \in B | \mu_t, 0 \leq t \leq t_0). \tag{5}$$

Now, by the finite definition of B, we have

$$1(Z^N(t_0) \in B) \to 1(Z(t_0) \in B)$$

and similarly for $\sigma Z^N(t_0)$, as for $N \geq N_\sigma$ we have

$$(\sigma Z(t_0))|_N = \sigma Z^N(t_0).$$

We can then prove Equation (5) by a standard limiting and monotone class argument. First, for any finite set of times $t_1, \ldots, t_n$ we have $\mu_{t_i}^N \to \mu_{t_i}$ for all $1 \leq i \leq n$ almost surely. Letting $f_i \colon P(S) \to \mathbb{R}$ be bounded and continuous, we then have

$$F^N = f_1(\mu_{t_1}^N) \cdots f_n(\mu_{t_n}^N) \to F = f_1(\mu_{t_1}) \cdots f_n(\mu_{t_n}),$$

almost surely. Now by 2.3 for $N \geq N_\sigma$ we have that

$$\mathbb{E}\left(1(\sigma Z^N(t_0) \in B)F^N\right) = \mathbb{E}\left(1(Z^N(t_0) \in B)F^N\right).$$

Applying the bounded convergence theorem, this shows that for any

$$F \in \sigma(\mu_t, 0 \leq t \leq t_0)$$

that can be factored as such, we have

$$\begin{aligned}
\mathbb{E}\left(1(\sigma Z(t_0) \in B)F\right) &= \mathbb{E}\lim_N \left(1(\sigma Z^N(t_0) \in B)F^N\right) \\
&= \lim_N \mathbb{E}\left(1\left(\sigma Z^N(t_0) \in B\right)F^N\right) \\
&= \lim_N \mathbb{E}\left(1\left(Z^N(t_0) \in B\right)F^N\right) \\
&= \mathbb{E}\left(1(Z(t_0) \in B)F\right).
\end{aligned}$$

By the Monotone Class Theorem (Theorem 5.1.5, [10]), the same equality holds for any bounded $F \in \sigma(\mu_t, 0 \leq t \leq t_0)$, proving Equation (5). As the finitely defined sets B are a basis for the topology on Π, the $\pi - \lambda$ theorem shows that Equation (5) holds for any measurable subset of Π, completing the proof. □

Having proved 3.3, we can now easily see that conditional on the past of the MC process, the partition at time t_0 is still given by the paintbox construction.

Lemma 3.4. *For any $k_1, k_2, \ldots, k_N \in \{1, 2 \ldots\}$, we have that*

$$\mathbb{P}(A^N = (k_1, \ldots, k_N)|\mu_t, 0 \leq t \leq t_0) = \mathbb{P}(A^N = (k_1, \ldots, k_N)|\mu_{t_0}).$$

Proof. 3.3 shows that conditional on $\mu_t, 0 \leq t \leq t_0$, the token partition $Z(t_0)$ is exchangeable and so given by a paintbox construction, which since the asymptotic rates (i.e mass partition) at time t_0 are determined by μ_{t_0} must coincide with the mass partition in the construction of μ_t. Therefore we have

$$\mathbb{P}(A^N = (k_1, \ldots, k_N)|\mu_t, 0 \leq t \leq t_0) = p_{k_1}(t_0) \cdots p_{k_n}(t_0)$$

which is clearly in $\sigma(\mu_{t_0})$. □

Next, we show that conditional on the past of the MC process, the future is determined solely by μ_{t_0} and the locations A^N of the first N tokens.

Lemma 3.5. *For $0 \leq t_0 \leq t_1$ we have that*

$$\mathbb{P}(\mu_{t_1}^N \in \cdot|\mu_t, 0 \leq t \leq t_0, A^N) = \mathbb{P}(\mu_{t_1}^N \in \cdot|\mu_{t_0}, A^N),$$

as random measures on $P(S)$.

Proof. The basic idea of this proof is that the locations of the first N tokens at time t_1 conditional on the past (before t_0) depends only on the locations of the first N tokens at time t_0.

For shorthand write $\xi_i(t)$ for the location in S of token i at time t, i.e.

$$\xi_i(t) = \xi_{u_i(t)} = s_{A_i},$$

and let $\xi^N(t)$ be the collection $\xi_i(t), 1 \le i \le N$ and similarly $\xi(t)$ for $\xi_i(t), i \ge 1$.

First, by construction the token process $\xi(t), t \ge 0$ is Markov and so as μ_t^N is a function of $\xi(t)$ we have

$$\begin{aligned}\mathbb{P}(\mu_{t_1}^N \in \cdot | \xi(t), 0 \le t \le t_0) &= \mathbb{P}(\mu_{t_1}^N \in \cdot | \xi(t_0)) \\ &= \mathbb{P}(\mu_{t_1}^N \in \cdot | \xi^N(t_0))\end{aligned}$$

using for the second equality that the path of the first N tokens doesn't depend on the path of the other tokens.

Therefore using that

$$\sigma(\mu_t, 0 \le t \le t_0, A^N) \subset \sigma(\xi(t), 0 \le t \le t_0)$$

we can complete the proof using the tower property of conditional expectation by calculating

$$\begin{aligned}&\mathbb{P}(\mu_{t_1}^N \in \cdot | \mu_t, 0 \le t \le t_0, A^N) \\ &\quad = \mathbb{E}\left(\mathbb{P}\left(\mu_{t_1}^N \in \cdot | \xi(t), 0 \le t \le t_0\right) | \mu_t, 0 \le t \le t_0, A^N\right) \\ &\quad = \mathbb{E}\left(\mathbb{P}\left(\mu_{t_1}^N \in \cdot | \xi^N(t_0)\right) | \mu_t, 0 \le t \le t_0, A^N\right) \\ &\quad = \mathbb{P}(\mu_{t_1}^N \in \cdot | \mu_{t_0}^N, A).\end{aligned}$$

For the last step, we use that

$$\sigma(\xi^N(t_0)) = \sigma(\mu_{t_0}^N, A^N) \subset \sigma(\mu_t, 0 \le t \le t_0, A^N).$$

□

We are now able to give our proof of 3.2.

Proof. Applying Lemma 3.4 and Lemma 3.5, by a simple calculation we have

$$\begin{aligned}&\mathbb{P}(\mu_{t_1}^N \in \cdot | \mu_t, 0 \le t \le t_0) \\ &\quad = \sum_{A^N = \bar{k}} \mathbb{P}(\mu_{t_1}^N \in \cdot, A^N = \bar{k} | \mu_t, 0 \le t \le t_0) \\ &\quad = \sum_{A^N = \bar{k}} \mathbb{P}(\mu_{t_1}^N \in \cdot | \mu_t, 0 \le t \le t_0, A^N)\, \mathbb{P}(A^N = \bar{k} | \mu_t, 0 \le t \le t_0) \\ &\quad = \sum_{A^N = \bar{k}} \mathbb{P}(\mu_{t_1}^N \in \cdot | \mu_{t_0}, A^N)\, \mathbb{P}(A^N = \bar{k} | \mu_{t_0}) \\ &\quad = \mathbb{P}(\mu_{t_1}^N \in \cdot | \mu_{t_1})\end{aligned}$$

where the sum is over all $\bar{k} = (k_1, \ldots, k_N) \subset \{1, 2, \ldots\}^N$. □

3.2 Proof of 3.1

We can now start the proof of 3.1 and show that $\mu_t, t \ge 0$ is Markov, by applying 3.2 and a limiting argument.

Proof. First, we claim that as random measures on $P(S)$ we have

$$\mathbb{P}(\mu_{t_1}^N \in \cdot | \mu_t, 0 \le t \le t_0) \to \mathbb{P}(\mu_{t_1} \in \cdot | \mu_t, 0 \le t \le t_0),$$

almost surely.

To see this, let $U \subset P(S)$ be open. Then as $\mu_{t_1}^N \to \mu_{t_1}$, if $\mu_{t_1} \in U$ then almost surely so is $\mu_{t_1}^N$ eventually; that is

$$\liminf_N 1(\mu_{t_1}^N \in U) \ge 1(\mu_{t_1} \in U).$$

Applying Fatou's Lemma we have that

$$\begin{aligned}\liminf_N \mathbb{E}(1(\mu_{t_1}^N \in U) | \mu_t, 0 \le t \le t_0) &\ge \mathbb{E}(\liminf_N 1(\mu_{t_1}^N \in U) | \mu_t, 0 \le t \le t_0) \\ &\ge \mathbb{E}(1(\mu_{t_1} \in U) | \mu_t, 0 \le t \le t_0),\end{aligned}$$

that is

$$\liminf_N \mathbb{P}(\mu_{t_1}^N \in U | \mu_t, 0 \le t \le t_0) \ge \mathbb{P}(\mu_{t_1} \in U | \mu_t, 0 \le t \le t_0)$$

almost surely. Applying the Portmanteau theorem to $P(S)$, we have proved the claim.

By the same argument, we also have

$$\mathbb{P}(\mu_{t_1}^N \in \cdot | \mu_{t_0}) \to \mathbb{P}(\mu_{t_1} \in \cdot | \mu_{t_0}),$$

almost surely as measures. By 3.2, we have that

$$\mathbb{P}(\mu_{t_1}^N \in \cdot | \mu_{t_0}) = \mathbb{P}(\mu_{t_1}^N \in \cdot | \mu_t, 0 \le t \le t_0)$$

almost surely, and so we can conclude that almost surely

$$\mathbb{P}(\mu_{t_1} \in \cdot | \mu_{t_0}) = \mathbb{P}(\mu_{t_1} \in \cdot | \mu_t, 0 \le t \le t_0)$$

completing the proof. □

All that is left to show to complete the proof of 3.1 is to prove that the MC process $\mu_t, t \ge 0$ is time-homogeneous. Write $\mathbb{P}_\mu$ for the distribution of the token process started at μ. Time homogeneity will follow immediately from the following lemma.

Lemma 3.6. *For any time $0 \le t_0$, the distribution of the token process conditional on $\mu_t, 0 \le t \le t_0$ is the same as that of the token process started from μ_{t_0}. That is*

$$\mathbb{P}_\mu(Z(t_0) \in \cdot | \mu_t, 0 \le t \le t_0) = \mathbb{P}_{\mu_{t_0}}(Z(0) \in \cdot).$$

Proof. This follows easily from 3.3, which shows that conditional on $\mu_t, 0 \le t \le t_0$, the location $A_i \in S$ of each token $i \ge 1$ is independent and given by the paintbox $p(t_0)$. It is immediate then that this is the same as the initial distribution of the token process for the countably supported measure μ_{t_0}, since the token process is Markov and both have the same initial distribution of token locations. □

4 Finite Support and Kingman's Coalescent

In this section, we begin our comparison between the token process construction and the Metric Coalescent, which we recall is originally only defined on $P_{\text{fs}}(S)$. Our main result here shows that on any compact set in S, the support of the token process $\mu_t, t \geq 0$ 'comes down from infinity'.

Proposition 4.1. *For any initial measure $\mu \in P(S)$ and any compact set $C \subset S$*

$$\#\left(\operatorname{supp}\mu_t \cap C\right) < \infty$$

for all $t > 0$, almost surely.

In fact, for any compact C and time $t > 0$, we can also show that the expected number of atoms in C at time t is finite.

Proposition 4.2. *For any initial measure $\mu \in P(S)$, any compact set $C \subset S$, and any $t > 0$ we have*

$$\mathbb{E}\left(\operatorname{supp}\mu_t \cap C\right) \leq \frac{2}{t\phi_{\min}(C)}.$$

Here we define $\phi_{\min}(C)$ to be the positive infinum of ϕ on C, that is

$$\phi_{\min}(C) = \inf_{x,y\in C} \phi(d(x,y)) > 0,$$

which by compactness of C is immediately bounded away from zero.

An easy consequence of 4.1 is our claim from Section 2.6 that the support of μ_t - as we have defined it, given by the set of token locations at time t - is nowhere dense and so matches the standard definition of support.

Lemma 4.3. *For all $t > 0$, the support of μ_t is nowhere dense and so almost surely*

$$\operatorname{supp}\mu_t = \overline{\{s \in S\colon \forall \mu_t(U) > 0 \text{ for all open } U \ni s\}}.$$

Also trivial, by the definition of the token process and Lemma 4.3, is that the support of μ_t is almost surely contained in that of the initial measure μ. This gives as an immediate corollary of 4.1 a portion of Item 2, namely that for any compactly supported initial measure $\mu \in P_{\text{cs}}(S)$ that $\mu_t, t \geq 0$ comes down from infinity.

Proposition 4.4. *If $\mu \in P_{\text{cs}}(S)$, then*

$$\#\operatorname{supp}\mu_t < \infty$$

for all $t > 0$, almost surely.

Note that while 4.4 shows that our extension of the MC comes down from infinity on $P_{\text{cs}}(S)$, this doesn't exclude it from doing so on a possibly larger class of measures in $P(S)$. We'll discuss this possibility in Section 10.

We begin with some bounds on $\#\left(\operatorname{supp}\mu_t^N \cap C\right)$ for $t > 0$. To do this, we compare the token process to the well studied Kingman Coalescent (KC), a classical coalescing partition process.

4.1 Kingman's Coalescent

Kingman's Coalescent is a classical coalescent process on the natural numbers $\mathbb{N}$ with applications throughout probability theory, in particular for population models in genetics where it satisfies an important universality property known as Möhle's Lemma. For our purposes, the connections between Kingman's Coalescent and the Metric Coalescent are deep, with the Metric Coalescent appearing as essentially a version of Kingman's model with dependent (exchangeable) rates.

Kingman's Coalescent is simple to describe, although a construction of it as a limit of finite processes takes some care. Simply put, it is a coalescing partition valued Markov process on $\mathbb{N}$, where each pair of blocks merges independently with some fixed rate $\gamma > 0$. Compare this to the MC, where to the token process, where the rates of merger between blocks are of course geometry dependent.

A standard fact about KC is that it *comes down from infinity*, that is, started from the initial partition of all singletons, for any positive time $t > 0$ the number of blocks $N_t^{KC} < \infty$ almost surely; in fact a stronger statement, bounding $\mathbb{E}\, N_t^{KC}$ holds as well [6].

While the rates of the token process are not independent, for any two atoms $s_i, s_j \in C$ of μ_t they satisfy

$$\nu_{ij} \geq \phi_{\min}(C) > 0,$$

for any compact set C. Thus by comparison to the KC we will show that the token process also comes down from infinity.

4.2 Coming Down From Infinity

While a direct comparison between the two process through a coupling may be possible, we find it simpler to apply the classical proof of finiteness for KC to the token process, with few adjustments. Our proof here follows [6] closely.

Fix some compact set $C \subset S$ and define

$$K_t(C) = \sum_{k=1}^{\infty} 1(u_k(t) = k, \xi_k \in C),$$

that is, K_t is the (possibly infinite) number of tokens that haven't merged by time t in C. Similarly we define

$$K_t^N = \sum_{k=1}^{N} 1(u_k(t) = k, \xi_k \in C),$$

for the N token process. Note that by our construction

$$K_t^N(C) = \#\left(\operatorname{supp} \mu_t^N \cap C\right)$$

as at most one alive token can occupy each point of S. Our main result in this section is the following.

Proposition 4.5. *For all $t > 0$, almost surely $K_t(C) < \infty$.*

We will follow the classical proof for Kingman's Coalescent to prove 4.5. As indicated, the exact same proof method borrowed from Kingman's Coalescent in fact yields the expectation bound in 4.2 for $K_t(C)$, (see Theorem 2.1 in [6]) whose details we omit here.

Note for 4.2 what will be important for our purposes is not the precise bound on $\mathbb{E}\, K_t(C)$ but only that such a bound exits for all $t > 0$. The rest this section will be outlining the proof of 4.5.

Consider a fixed initial measure $\mu \in P(S)$. For $1 \leq n \leq N$, let

$$T_n^N = \inf\{t \geq 0 \colon K_t^N(C) = n\}.$$

The key to our proof is that while the pure death process $K_t^N(C)$ is typically not Markovian, the token process is and so when there are $K_t^N(C) = n$ blocks left, the conditional intensity of a transition from n to $n-1$ is at least $\binom{n}{2}\phi_{\min}(\mu)$, the rate at which two atoms of μ_t^N within C merge. In fact, it is a priori much higher than this, since an atom from C can merge into an atom outside of C.

We will make use of the following estimate of T_n^N.

Lemma 4.6. *For all* $1 \leq n \leq N$

$$\mathbb{E}\, T_n^N \leq \frac{2}{\phi_{\min}(C)n}.$$

Proof. At time T_n^N let $k_1, \ldots, k_n$ be the tokens still alive in C. As the N token process is Markov, from the memoryless property of exponential random variables we have that conditional on the process up to time T_n^N, the later meeting times $t_{k_i,k_j} - T_n^N$ are distributed as independent exponentials of rate ν_{k_i,k_j}. Therefore, conditional on the past, the next meeting $T_{n-1}^N - T_n^N$ is distributed as an exponential of rate

$$\theta \geq \sum_{i \neq j \leq n} \nu_{k_i,k_j},$$

where inequality comes from the rate at which the atoms $k_1, \ldots, k_n$ merge with atoms outside of C that are still alive.

As $\nu_{a,b} \geq \phi_{\min}(C)$ for all $a, b \in C$, we can conclude that $\theta \geq \binom{n}{2}\phi_{\min}(C)$.

This allows us to calculate that

$$\mathbb{E}\, T_{n-1}^N - T_n^N \leq \frac{2}{\phi_{\min}(C)n(n-1)},$$

and so

$$\begin{aligned}
\mathbb{E}\, T_n^N &= \mathbb{E} \sum_{m=n-1}^{N} T_m^N - T_{m+1}^N \\
&\leq \sum_{m=n}^{N-1} \frac{2}{\phi_{\min}(C)m(m+1)} \\
&= \frac{2}{\phi_{\min}(C)} \left(\frac{1}{n} - \frac{1}{N}\right).
\end{aligned}$$

□

From this we can easily complete our proof of 4.5.

Proof. Applying Markov's inequality to Lemma 4.6, we find that for any $n \geq 1$

$$\mathbb{P}(K_t^N(C) \geq n) = \mathbb{P}(T_n^N \geq t) \leq \frac{2}{\phi_{\min}(C)tn}.$$

As $K_t^N \uparrow K_t$ almost surely, this implies that

$$\mathbb{P}(K_t(C) \geq n) \leq \frac{2}{\phi_{\min}(C)tn}$$

which easily gives that $K_t(C) < \infty$ almost surely, completing the proof. □

The proof of 4.2 follows by a similar argument, looking at second moment estimates of the stopping times T_n^N (see Section 2.1.2, [6]).

Applying a standard countable additivity argument to 4.5 at a descending sequence of times $t_i \downarrow 0$ we conclude our proof of 4.1.

5 The Metric Coalescent

In this section, we justify calling our constructed process an extension of the Metric Coalescent, by showing that from any finitely supported initial measure $\mu \in P_{\text{fs}}(S)$, the process $\mu_t, t \geq 0$ is distributed as the MC as defined in Section 1.2. Our main result in this section is the following.

Proposition 5.1. *If $\mu \in P_{fs}(S)$, then the process $\mu_t, t \geq 0$ starting from $\mu_0 = \mu$ is distributed as the Metric Coalescent.*

In Section 4 we've shown that for an initial measure $\mu \in P_{\text{cs}}(S)$, for all times $t > 0$ the process is contained in $P_{\text{fs}}(S)$. Combined with the Markov property, 5.1 proves the claim in Item 2 that if $\mu \in P_{\text{cs}}(S)$, then for all $t_0 > 0$, the process $\mu_t, t \geq t_0$ is also distributed as the Metric Coalescent.

5.1 The Partial Order $\prec$

The evolution of the MC process $\mu_t, t \geq 0$ is defined in terms of the asymmetric description given by the token process, where tokens merge into lower tokens. Viewing the process through the symmetric description

$$\mu_t = \sum_{k=1}^{\infty} p_k(t)\delta(s_k),$$

we no longer know the locations of the tokens and so no longer can determine the future of the process deterministically from the meeting times.

From 3.3 however, we know how the distribution of the tokens at time t depends upon μ_t and so know the relative rankings of the lowest elements of the partitions at each atom s_k, or equivalently the blocks of the paintbox partition. This is known as the size biased permutation of the paintbox partition.[13]

Conceptually we find it simpler to view the size biased permutation as a (random) partial ordering $\prec$ on S. Let $\mathfrak{O}(S)$ be the set of all partial orders on

S. The order $\prec$ can be defined easily from the initial locations of the tokens $\xi_i, i \geq 1$. For any $s \in S$, write $l(s)$ for the lowest token starting at s, i.e.

$$l(s) = \inf\{i \geq 1 \colon \xi_i = s\}$$

with the convention that $l(s) = \infty$ if no such token exists. Then for $s, \tilde{s} \in S$ we define $s \prec \tilde{s}$ if

$$l(s) < l(\tilde{s}) < \infty.$$

Note that for all $t > 0$, by Lemma 4.3

$$\operatorname{supp} \mu_t \subset \{\xi_1, \xi_2, \ldots\}$$

and so we always have that $\prec$ is a full linear order when restricted to $\operatorname{supp} \mu_t$.

Given the initial locations $\xi_i, i \geq 1$ the ordering $\prec$ is fully known. We however are interested in the conditional distribution of $\prec$ given the starting measure $\mu_0 = \mu$. The key motivation here, is that when μ_0 is finitely supported and so meeting rates are discrete, the future distribution of the process will be seen to be determined entirely by just the locations of the atoms of μ_t and the ordering $\prec$ of the tokens they contain.

5.2 Size Biased Ordering

Our first goal in this section is to explicitly calculate the conditional distribution of $\prec$ on $\operatorname{supp} \mu_{t_0}$ at some time $t_0 \geq 0$ given $\mu_t, 0 \leq t \leq t_0$. We will show that this distribution is given by the size biased ordering, to be defined shortly.

Proposition 5.2. *The conditional distribution of $\prec$ given the process $\mu_t, 0 \leq t \leq t_0$ is size biased ordering. That is, for any $s_1, \ldots, s_n \in \operatorname{supp} \mu_{t_0}$*

$$\mathbb{P}(s_1 \prec \ldots \prec s_n | \mu_t, 0 \leq t \leq t_0) = \mathrm{SBO}(\mu_{t_0}, s_1, \ldots, s_n).$$

5.2.1 Definition of Size Biased Ordering

For any $s_1, \ldots, s_n \in S$ and $\nu \in P(S)$ write $\mathrm{SBO}(\nu, s_1, \ldots, s_n)$ as shorthand for the formula

$$SBO(\nu, s_1, \ldots, s_n) = \mathbb{P}(\mathrm{Exp}(\nu(s_n)) < \ldots < \mathrm{Exp}(\nu(s_1))), \tag{6}$$

where here $\mathrm{Exp}(\mu(s_n))$ are independent exponential random variables of rate $\mu(s_n)$. This is one of several equivalent definitions of the size-biased random ordering [13] associated with ν. Explicitly this can also be written as

$$SBO(\nu, s_1, \ldots, s_n) = \frac{1(\nu(s_i) \neq 0, 1 \leq i \leq n)\nu(s_1) \cdots \nu(s_n)}{(\nu(s_1) + \ldots + \nu(s_n))(\nu(s_2) + \ldots) \cdots \nu(s_n)},$$

but we conceptually prefer the first definition.

We say that a random partial order $\tilde{\prec} \in \mathfrak{O}(S)$ has **ν-size biased ordering** if for any $s_1, \ldots, s_n \in S$

$$\mathbb{P}(s_1 \tilde{\prec} \ldots \tilde{\prec} s_n) = SBO(\nu, s_1, \ldots, s_n)$$

To see that conditionally $\prec$ has size biased ordering, proving 5.2, we need only recall 3.3 which gives an explicit formula for the conditional location of the tokens. That the ordering then satisfies Equation (6) is then standard. [13]

5.3 Proof of 5.1

We are now able to give a short proof of 5.1.

Proof. We will prove this inductively over the jumps T_i of $\mu_t, t \geq 0$.

Since $\mu_0 \in P_{\text{fs}}(S)$, each jump clearly occurs when two atoms of μ_t merge. The meeting rates between two atoms s_i and s_j are given by the meeting rates of the lowest tokens at s_i and s_j - however the rate is $\phi(d(s_i, s_j))$ which clearly depends only on s_i and s_j and not on whatever tokens are located there. This meeting process is obviously Markovian as the meeting rates are exponentially distributed (and so memoryless) and clearly matches the description of the meeting rates of the Metric Coalescent in Section 1.2.

All that's left is to then check the distribution of which token merges with which. Consider the first meeting time T_1, between two atoms s_{i_1} and s_{j_1}. Which token absorbs the other depends only on their relative ranking, i.e. whether $s_{i_1} \prec s_{j_1}$ or vice versa. By 5.2 this is given by

$$\mathbb{P}(s_{i_1} \prec s_{j_1})|\mu_0) = \frac{\mu_0(s_{i_1})}{\mu_0(s_{i_1}) + \mu_0(s_{j_1})}$$

matching that of the Metric Coalescent.

Similarly, writing the first n jump (i.e. meeting) times as $T_0 = 0, T_1, \ldots, T_n$, if at T_n two atoms s_{i_n} and s_{j_n} meet, then the probability of s_{i_n} absorbing s_{j_n} is given by

$$\begin{aligned}\mathbb{P}(s_{i_n} \prec s_{j_n}|\mu_{T_{n-1}}, \ldots, \mu_{T_0}) &= SBO(\mu_{T_{n-1}}, s_{i_n}, s_{j_n}) \\ &= \frac{\mu_{T_{n-1}}(s_{i_n})}{\mu_{T_{n-1}}(s_{i_n}) + \mu_{T_{n-1}}(s_{j_n})}\end{aligned}$$

from 5.2. This again clearly matches the Metric Coalescent proving that the two finite jump processes are equal in distribution. □

6 Martingales of the Metric coalescent

In this section, we analyse the process $\mu_t, t \geq 0$ by studying a family of associated martingales. In particular, we use martingale arguments to study some of the path properties of our construction. Our main result in this section is the following.

Proposition 6.1. *From any initial measure $\mu_0 = \mu$ the process μ_t is right continuous at $t = 0$ almost surely.*

In particular, the same argument (or the Markov property) can be used to show that for any fixed time t_0, the MC process $\mu_t, t \geq 0$ is continuous at t_0 almost surely. However our proof later that the MC process is cadlag almost surely (see Section 8) only requires right-continuity at time $t = 0$, so we'll focus on that here.

Of course, the process $\mu_t, t \geq 0$ is - for generic initial $\mu \in P(S)$ - not path continuous almost surely. This can easily be seen since from any initial measure in $P_{\text{fs}}(S)$ the process is a discrete jump process and so is never path continuous, only cadlag. For now we'll defer consideration of the path properties of the process until after we prove Feller continuity (see Section 7).

To prove 6.1, we consider the evolution of functions integrated against our measure-valued process. Specifically, for any $f \in C_b(S)$ we consider the process $\mu_t(f), t \geq 0$ and equivalently $\mu_t^N(f), t \geq 0$ for the N-th empirical measures.

Our main result in the study of these processes is that these processes are martingales with explicitly computable quadratic variation.

Proposition 6.2. *For any $f \in C_b(S)$, $\mu_t(f), t \geq 0$ and $\mu_t^N(f), t \geq 0$ are martingales with quadratic variation satisfying*

$$\mathbb{E}\left(\mu_t(f) - \mu_0(f)\right)^2 = \frac{1}{2}\mathbb{E}\left(1 - \exp(-\phi(d(\xi_1, \xi_2))t)\right)\left(f(\xi_1) - f(\xi_2)\right)^2$$

for any time $t \geq 0$.

We begin by studying the process dynamics of $\mu_t^N(f)$ and use that to prove the first half of 6.2. Then, a study of the quadratic variation for the case of countably supported initial measures will lead the result for general initial measures.

6.1 Process Dynamics of the Metric Coalescent

We begin by showing that the processes $\mu_t^N(f), t \geq 0$ are martingales. We will need to make use of the conditional size biased permutation of the partitions of μ_t^N at any time $t > 0$, as seen in Lemma 6.3. This simple extension of 5.2 follows easily from the fact that the empirical measures μ_t^N also as evolve as the MC, in this case as the originally defined MC from the (random) initial measure μ_0^N.

Write

$$\mu_t^N = \sum_{i=1}^{K^N(t)} p_i \delta(s_i),$$

where $K^N(t)$ is the finite number of atoms of μ_t^N at time t.

Lemma 6.3. *At any fixed time $t \geq 0$, the size biased permutation of μ_t^N satisfies*

$$\mathbb{P}(s_i \prec s_j | \mu_t^N) = \frac{p_i}{p_i + p_j}.$$

Let f be any bounded continuous function on S.

Lemma 6.4. *$\mu_t^N(f), t \geq 0$ is a martingale.*

Proof. We show this from the process dynamics of $\mu_t^N(f)$. For some fixed time $t \geq 0$ write

$$\mu_t^N = \sum_{i=1}^{K^N(t)} p_i \delta(s_i),$$

with the support of μ_t^N in any order. For simplicity of notation write

$$\nu_{ij} = \phi(d(s_i, s_j))$$

for the rate of coalescence between the atoms s_i and s_j. Then

$$\begin{aligned}\mathbb{E}\left(d\mu_t^N(f)|\mu_s, 0 \le s \le t\right) \\ &= \sum_{i\neq j} \nu_{ij}\Big(\frac{p_i}{p_i+p_j}p_j(f(s_i)-f(s_j)) + \frac{p_j}{p_i+p_j}p_i(f(s_j)-f(s_i))\Big)\,dt \\ &= 0\,dt\end{aligned}$$

since from time t, at rate ν_{ij} the atoms at s_i and s_j merge with - by Lemma 6.3 - s_i winning the meeting (i.e. having the lower token) and absorbing the mass p_j of s_j with probability $\frac{p_i}{p_i+p_j}$ and similarly for s_j winning. □

As the sequence of martingales $\mu_t^N(f), t \ge 0$ are uniformly bounded since f is, standard convergence results imply that $\mu_t(f), t \ge 0$ is itself a martingale [15]. Conversely, as the sequence is uniformly integrable this can be shown directly without much difficulty.

To prove the second claim in 6.2 for μ_t^N we begin by analysing the second moment of $\mu_t(f)$ for a countably supported initial measure μ.

6.2 Second Moment Calculation

Consider a fixed countably supported initial measure

$$\mu_0 = \sum_{i=1}^{\infty} p_i(0)\delta(s_i).$$

By construction, for all times $t \ge 0$ the atoms of μ_t are contained within those of μ_0 and so we may write

$$\mu_t = \sum_{i=1}^{\infty} p_i(t)\delta(s_i),$$

with $p_i(t)$ random and possibly zero. For each $i \ge 1$, the process $p_i(t), t \ge 0$ is a martingale as can be seen easily (when $\{s_i\}_{i\ge1}$ is nowhere dense, otherwise by the same argument as for the higher moments) choosing some $f \in C_b(S)$ separating s_i from the other atoms of μ_0. Therefore the expected mass at s_i at time t is given by

$$\mathbb{E}\,p_i(t) = p_i(0),$$

for all $t \ge 0$. We begin by calculating the second moments of these masses.

Proposition 6.5. *For any $i \neq j \ge 1$ and $t \ge 0$ we have*

$$\mathbb{E}\,p_i(t)p_j(t) = p_i(0)p_j(0)\exp(-\phi(d(s_i,s_j))t).$$

Similarly we have

$$\mathbb{E}\,p_i^2(t) = \sum_{j\ge1} p_i(0)p_j(0)\left(1-\exp(-\phi(d(s_i,s_j))t)\right).$$

Proof. First, by the paintbox construction of the token process at time t, we have that $p_i(t)p_j(t)$ is the probability of finding any two fixed tokens at s_i and s_j respectively. Thus

$$\mathbb{E}\, p_i(t)p_j(t) = \mathbb{P}(\text{ token 1 is at } s_i, \text{ token 2 is at } s_j \text{ at time } t).$$

Now, the only way this can occur is if $\xi_1 = s_i, \xi_2 = s_j$ and tokens 1 and 2 haven't met by time t, that is $t_{12} \geq t$. By the definition of the token process (or equivalently its paintbox construction at time zero) we can conclude that

$$\begin{aligned}\mathbb{E}\, p_i(t)p_j(t) &= \mathbb{P}(\xi_1 = s_i, \xi_2 = s_j, t_{12} \geq t) \\ &= p_i(0)p_j(0)\, \mathbb{P}(t_{12} \geq t | \xi_1 = s_i, \xi_2 = s_j) \\ &= p_i(0)p_j(0) \exp(-\phi(d(s_i, s_j))t).\end{aligned}$$

The second statement follows by a similar proof, noting that if both tokens 1 and 2 are at s_i at time t, then token 1 must have started at s_i and their first meeting must satisfy $t_{12} \leq t$. □

In theory this same idea gives explicit formulas for higher moments of $p_i(t), i \geq 1$, where an n-th moment corresponds to a certain subset of meeting trees on n tokens. In practice however, for anything past the second moment the number of possible meeting trees makes such a calculation impractical. We discuss this further in Section 7.5.

Using 6.5 we are able to calculate the quadratic variation of $\mu_t(f)$ from a countably supported initial measure μ_0.

Lemma 6.6. *For any countably supported initial measure*

$$\mu_0 = \sum_{i=1}^{\infty} p_i(0)\delta(s_i)$$

and $f \in C_b(S)$, the quadratic variation of $\mu_t(f)$ satisfies

$$\mathbb{E}\left(\mu_t(f) - \mu_0(f)\right)^2 = \sum_{i \neq j} p_i(0)p_j(0)\left(1 - \exp(-\phi(d(s_i, s_j))t)\right)\left(f(s_i) - f(s_j)\right)^2.$$

Proof. This follows by an application of 6.5 to

$$\mu_t(f) = \sum_{i \geq 1} p_i(t)f_i(t),$$

after some simplification, recalling that $\mu_t(f), t \geq 0$ is a martingale and so

$$\mathbb{E}\left(\mu_t(f) - \mu_0(f)\right)^2 = \mathbb{E}\left(\mu_t(f)\right)^2 - \left(\mu_0(f)\right)^2.$$

□

We can now complete our proof of 6.2 by applying Lemma 6.6 to the sequence of empirical measures $\mu_t^N, N \geq 1$ - which for any initial measure $\mu \in P(S)$ are still countably supported.

Proof. Write $F_t \colon S^2 \to \mathbb{R}$ for the function

$$F_t(x,y) = \begin{cases} (1-\exp(-\phi(d(x,y))t))\,(f(x)-f(y))^2 & \text{if } x \neq y \\ 0 & \text{if } x = y. \end{cases}$$

Importantly note that $F_t(x,y)$ is continuous, in particular along the diagonal $x = y$. This is easy to see as $(1-\exp(-\phi(d(x,y))t))$ is bounded and so if for some sequences $x_i, i \geq 1$ and $y_i, i \geq 1$ we have $d(x_i, y_i) \to 0$ then

$$(f(x_i) - f(y_i))^2 \to 0$$

showing that $F_t(x_i, y_i) \to 0$. Since by assumption $f \in C_b(S)$ is bounded, so is $F_t(x,y)$ and so clearly $\mathbb{E}\, F_t(\xi_1, \xi_2)$ is well defined and also bounded.

We claim that

$$\mathbb{E}\left(\left(\mu_t^N(f) - \mu_0^N(f)\right)^2 |\mu_0^N\right) = \sum_{i,j=1}^{N} \frac{1}{2N^2} F_t(\xi_i, \xi_j). \tag{7}$$

Lemma 6.6 gives us a formula for $\mathbb{E}\left(\left(\mu_t^N - \mu_0^N\right)^2 |\mu_0^N\right)$, however some care is needed to prove Equation (7) as

$$\mu_0^N = \frac{1}{N}\sum_{i=1}^{N} \delta(\xi_i)$$

may not be a atomic decomposition of μ_0^N as some tokens may already be at the same location at time $t = 0$. To see that Equation (7) still holds, consider two atoms $p_1\delta(s_1)$ and $p_2\delta(s_2)$ of μ_0^N. Then we have that

$$\begin{aligned} p_1 p_2 F_t(s_1, s_2) &= \left(\sum_{i=1}^{N} \frac{1(\xi_i = s_1)}{N}\right)\left(\sum_{i=1}^{N} \frac{1(\xi_j = s_2)}{N}\right) F_t(s_1, s_2) \\ &= \sum_{i=1}^{N}\sum_{j=1}^{N} \frac{1(\xi_i = s_1)1(\xi_j = s_2)}{N^2} F_t(s_1, s_2) \\ &= \sum_{i=1}^{N}\sum_{j=1}^{N} \frac{1(\xi_i = s_1)1(\xi_j = s_2)}{N^2} F_t(\xi_i, \xi_j). \end{aligned}$$

Thus writing an atomic decomposition for μ_0^N as

$$\mu_0^N = \sum_{k=1}^{K} p_k \delta(s_k)$$

we have that, summing over all pairs of atoms of μ_0^N and using that F_t is zero on the diagonal, i.e

$$F_t(s_k, s_k) = 0$$

for all $1 \leq k \leq K$ we have that

$$\begin{aligned}\mathbb{E}\left(\left(\mu_t^N(f) - \mu_0^N(f)\right)^2 |\mu_0^N\right) &= \sum_{k_1 \neq k_2} p_{k_1} p_{k_2} F_t(s_{k_1}, s_{k_2}) \\ &= \frac{1}{2} \sum_{k_1=1}^{K} \sum_{k_2=1}^{K} p_{k_1} p_{k_2} F_t(s_{k_1}, s_{k_2}) \\ &= \frac{1}{2} \sum_{k_1=1}^{K} \sum_{k_2=1}^{K} \sum_{i=1}^{N} \sum_{j=1}^{N} \frac{1(\xi_i = s_{k_1}) 1(\xi_j = s_{k_2})}{N^2} F_t(\xi_i, \xi_j) \\ &= \frac{1}{2} \sum_{i=1}^{N} \sum_{j=1}^{N} \frac{1}{N^2} F_t(\xi_i, \xi_j),\end{aligned}$$

proving Equation (7).

To finish the proof of 6.2, taking the expectation of Equation (7) using the exchangeability of the initial locations $\xi_i, 1 \leq i \leq N$ we have that

$$\begin{aligned}\mathbb{E}\left(\mu_t^N(f) - \mu_0^N(f)\right)^2 &= \frac{N^2}{2N^2} \mathbb{E}\, F_t(\xi_1, \xi_2) \\ &= \frac{1}{2} \mathbb{E}\, F_t(\xi_1, \xi_2).\end{aligned}$$

By 2.1, $\mu_t^N \to \mu_t$ weakly almost surely and so $\mu_t^N(f) \to \mu_t(f)$ almost surely; and similarly for μ_0^N. Therefore as f is bounded we can apply the Bounded Convergence theorem to complete the proof. □

We are now able to complete our proof of 6.1.

6.3 Proof of 6.1

To show the almost sure continuity of $\mu_t, t \geq 0$ at time $t = 0$, we first show the almost sure continuity of $\mu_t(f)$ at $t = 0$ for any $f \in C_b(S)$ and then apply 2.10 as was done in the proof of 2.1.

Lemma 6.7. *For a fixed function $f \in C_b(S)$ the process $\mu_t(f), t \geq 0$ is almost surely continuous at $t = 0$.*

Proof. First, by 6.2 $\mu_t(f), t \geq 0$ is a bounded martingale. So, applying the martingale convergence theorem for backwards martingales - Theorem 4.6.1, [10] - at $t = 0$ we see that $\mu_t(f), t \geq 0$ has an almost sure limit $M_{\lim}^f$ as $t \downarrow 0$. So, to complete the proof we need to show that $M_{\lim}^f = \mu_0(f)$ almost surely.

The second claim of 6.2 tells us that

$$\mathbb{E}\left(\mu_t(f) - \mu_0(f)\right)^2 = \mathbb{E}\, F_t(\xi_1, \xi_2)$$

writing $F_t(x, y)$ as in the proof of 6.2. We claim that

$$F_t(\xi_1, \xi_2) \to 0$$

almost surely as $t \downarrow 0$. To see this, when $\xi_1 = \xi_2$ already we have $F_t(\xi_1, \xi_2) = 0$ as

$$f(\xi_1) = f(\xi_2).$$

When $\xi_1 \neq \xi_2$, the limit follows from

$$1 - \exp(-\phi(d(\xi_1, \xi_2))t) \to 0$$

and the fact that f is bounded. Thus applying the Bounded Convergence theorem to $F_t(\xi_1, \xi_2)$ we get that

$$\mathbb{E}\left(\mu_t(f) - \mu_0(f)\right)^2 \to 0,$$

as $t \downarrow 0$ showing that $\mu_t(f) \to \mu_0(f)$ in L^2. As we already know that $\mu_t(f), t \geq 0$ has an almost surely limit $M^f_{\lim}$ at $t = 0$, we can conclude that $\mu_0(f) = M^f_{\lim}$ almost surely completing the proof. □

To conclude that $\mu_t, t \geq 0$ is right continuous at $t = 0$, we need only recall the standard Theorem 2.9 giving a countable convergence determining class for weak convergence. Therefore, applying 2.10 with Lemma 6.7 we complete the proof of 6.1.

7 Feller Continuity

The final step of our proof of existence in Item 2 is to show that our extended Metric Coalescent process μ_t is Feller continuous. Following the terminology of [17], write $C_b(P(S))$ for continuous bounded functions on $P(S)$. Then to show that our process satisfies the definition of "Feller continuous" we must show that for $F \in C_b(P(S))$ and any $t \geq 0$, the transition operator P_t on $C_b(P(S))$ given by

$$(P_t F)(x)\, \mathbb{E}_{\mu_0 = x}\, F(\mu_t)$$

is also in $C_b(P(S))$. In our setting, this is equivalent to showing that the random measures μ_t are weakly continuous, that is, if $\mu^{(i)} \to \mu$ weakly in $P(S)$, then

$$\mu_t^{(i)} \to \mu_t^{(i)}$$

weakly in $P(P(S))$. Our goal in this section is to prove Feller continuity.

Proposition 7.1. *The extension of the Metric coalescent process is Feller continuous.*

Note the slight distinction here between a Feller continuous process and the other definition (e.g. [17]) of a "Feller process" i.e. a process generated by a Feller semi-group $P_t, t \geq 0$, whose state space is required to be locally compact, which (generically) does not hold in our setting for $P(S)$. Nonetheless, we will still show that a few properties of Feller processes still hold, in particular the existence of a cadlag version – in our case the constructed version, see 8.1 – and its subsequent uniqueness in distribution.

In the case when S is compact, the space of Borel probability measures $P(S)$ is compact (and so trivially locally compact) and as an immediate consequence of 7.1 the MC can be seen to be a Feller process.

7.1 Convergence of Random Measures

Our goal in this section is to show the weak convergence of the MC processes for all times $t \geq 0$ for a sequence of weakly convergent initial measures.

Proposition 7.2. *Let $\mu^{(i)} \to \mu$ in $P(S)$ and $t \geq 0$ any fixed time. Then*

$$\mu_t^{(i)} \to \mu_t$$

weakly in $P(P(S))$.

For time $t = 0$, this is trivial, so we focus on proving this for positive times t. Our proof relies on the following criteria for weak convergence of random measures ([9], Theorem 11.1.8), as well as a standard moment method for convergence in distribution of bounded random variables (Section 2.3.e, [10]).

Theorem 7.3. *Let $\nu^i, 1 \leq i \leq \infty$ be a sequence of random measures in $P(S)$. Then the following are equivalent:*

1. *$\nu^i \to \nu^\infty$ weakly in $P(P(S))$.*
2. *For all $f \in C_b(S)$ with bounded support,*

$$\nu^i(f) \xrightarrow{d} \nu^\infty(f).$$

3. *For all $f \in C_b(S)$ with bounded support and $k \geq 1$*

$$\mathbb{E}(\nu^i(f))^k \to \mathbb{E}(\nu^\infty(f))^k.$$

The outline of our proof of 7.2 is as follows: for any fixed function $f \in C_b(S)$, we first reduce the moment calculation to a compactly supported approximation of f, and then reduce the comparison between $\mu_t^{(t)}$ and μ_t to a comparison between their N-th token processes. To do this, we give a coupling of the token processes. The key proposition in the proof is the following.

Proposition 7.4. *Let $\mu^{(i)} \to \mu$ weakly in $P(S)$. Then for all $t \geq 0, N \geq 1$ and $f \in C_b(S)$, there is an $I(N, f)$ such that if $i \geq I$ then*

$$\inf_\Gamma \mathbb{E}_\Gamma(\mu_t^{(i),N}(f) - \mu_t^N(f))^2 \leq \frac{f_{\max} + 1}{N},$$

where the infinum is over all couplings Γ of $\mu_t^{(i),N}$ and μ_t^N.

In fact, our proof of 7.4 shows that as $i \to \infty$

$$\inf_\Gamma \mathbb{E}_\Gamma(\mu_t^{(i),N}(f) - \mu_t^N(f))^2 \to 0,$$

however for notational purposes it is far more convenient (and still sufficient) to leave our result as stated. See Section 7.5 for discussion of an alternative approach to 7.4.

A quick note on notation: we will use $\mu^{(i)}$ to refer to the sequence of initial measures and $\mu_t^{(i)}, t \geq 0$ to the associated processes. The corresponding N-th token process will be written as $\mu_t^{(i),N}$ for $\mu^{(i)}$ and as usual μ_t^N for μ.

7.2 Proof of 7.2

We defer the proof 7.4 which is rather technical and first complete the proof of 7.2. We will need the following convergence bound.

Lemma 7.5. *Let $f \in C_b(S)$ be bounded by $|f| \leq f_{\max}$. Then for any $\nu \in P(S)$ and all $t \geq 0$*

$$\mathbb{E}\,|\nu_t^N(f) - \nu_t(f)| \leq \frac{f_{\max}}{\sqrt{N}}.$$

Proof. This is essentially an extension of the almost sure convergence shown in Lemma 2.8, using boundedness to prove convergence in L^2. Write

$$\nu_t^N(f) = \frac{1}{N}\sum_{i=1}^{N} f_i(t)$$

where $f_i(t)$ are conditionally independent given ν_t, defined by

$$f_i(t) = \sum_{i=k}^{\infty} f(s_k(t))1(i \in B_k(t)),$$

and thus bounded by $|f_i(t)| \leq f_{\max}$. Recall from the paintbox partition that B_k is the block of $Z(t)$ corresponding to the atom s_k of ν_t.

Then, by conditional independence we have

$$\mathbb{E}\left(\frac{1}{N}\sum_{i=1}^{N} f_i(t) - \nu_t(f)\right)^2 \leq \frac{f_{\max}^2}{N},$$

from which our claim follows easily. □

Our general method for proving 7.2 will be to first reduce the calculation for f to a compactly supported approximation f_ϵ, then by the inherent coupling of the token process to the MC reduce to a calculation for $\mu_t^{(i),N}$ and μ_t^N. Finally, applying 7.4 the result will follow easily.

We make use of the following lemma on compactly supported approximation.

Lemma 7.6. *For any $f \in C_b(S)$ and compact $K \subset S$, there exists a compactly supported $\tilde{f} \in C_0(S)$ such that:*

1. *$|\tilde{f}(x)| \leq |f(x)|$ for all $x \in S$,*
2. *$\tilde{f}(x) = f(x)$ for all $x \in K$.*

This is essentially an application of Urysohn's Lemma (Theorem 2.12, [1]).

We are now ready to prove 7.2.

Proof. By Item 3, it suffices to consider a fixed $f \in C_b(S)$ and $k \geq 1$. We then need to show that

$$\mathbb{E}\left(\mu_t^{(i)}(f)\right)^k \to \mathbb{E}\left(\mu_t(f)\right)^k.$$

Write $f_{\max}$ for the supremum of $|f|$ on S.

First, by assumption $\mu^{(i)} \to \mu$ and so the measures in $P(S)$ are tight. Thus, for any $\epsilon > 0$ there is a compact $K_\epsilon \subset S$ such that

$$\mu^{(i)}(K_\epsilon) \geq 1 - \epsilon$$

and similarly for μ. Let f_ϵ be the continuous, compactly supported extension of f on K_ϵ guaranteed by Item 2.

We claim that

$$\mathbb{E}\,|\mu_t(f) - \mu_t(f_\epsilon)| \leq 2f_{\max}\epsilon, \tag{8}$$

and similarly for $\mu_t^{(i)}$. This can be seen easily from

$$\begin{aligned}
\mathbb{E}\,|\mu_t(f) - \mu_t(f_\epsilon)| &\leq \mathbb{E}\,\mu_t(|f - f_\epsilon|) \\
&\leq \mathbb{E}\,\mu_0(|f - f_\epsilon|) \\
&= 2f_{\max}\,\mathbb{E}\,\mu_0(S \setminus K_\epsilon) \\
&\leq 2f_{\max}\epsilon
\end{aligned}$$

using the fact that $\mu_t(|f - f_\epsilon|)$ is a bounded martingale (as in 6.2).

Next, we note that for any X, Y with $|X|, |Y| \leq f_{\max}$ that

$$|X^k - Y^k| \leq k f_{\max}^{k-1}|X - Y|$$

by the simple Lipschitz bound on $f(x) = x^k$.

Putting this all together, we can bound

$$|E(\mu_t(f))^{2k} - \mathbb{E}(\mu_t^{(i)}(f))^{2k}|$$

as follows:

$$\begin{aligned}
&|\,\mathbb{E}(\mu_t(f))^k - \mathbb{E}(\mu_t^{(i)}(f))^k| \\
&\quad\leq |\,\mathbb{E}(\mu_t(f))^k - \mathbb{E}(\mu_t(f_\epsilon))^k| + |\,\mathbb{E}(\mu_t(f_\epsilon))^k - \mathbb{E}(\mu_t^{(i)}(f_\epsilon))^k| + |\,\mathbb{E}(\mu_t^i(f_\epsilon))^k - \mathbb{E}(\mu_t^i(f))^k| \\
&\quad\leq \mathbb{E}\,|(\mu_t(f))^k - (\mu_t(f_\epsilon))^k| + \mathbb{E}\,|(\mu_t^{(i)}(f)^k - (\mu_t^{(i)}(f_\epsilon))^k| + |\,\mathbb{E}(\mu_t(f_\epsilon))^k - \mathbb{E}(\mu_t^{(i)}(f_\epsilon))^k| \\
&\quad\leq k f_{\max}^{k-1}\,\mathbb{E}\left(|\mu_t(f) - \mu_t(f_\epsilon)| + |\mu_t^{(i)}(f) - \mu_t^{(i)}(f_\epsilon)|\right) + |\,\mathbb{E}(\mu_t(f_\epsilon))^k - \mathbb{E}(\mu_t^{(i)}(f_\epsilon))^k| \\
&\quad\leq 4k f_{\max}^k\epsilon + \mathbb{E}\,|(\mu_t(f_\epsilon))^k - (\mu_t^N(f_\epsilon))^k| + \mathbb{E}\,|(\mu_t^{(i)}(f_\epsilon))^k - (\mu_t^{(i),N}(f_\epsilon))^k| \\
&\quad+ |\,\mathbb{E}(\mu_t^N(f_\epsilon))^k - \mathbb{E}(\mu_t^{(i),N}(f_\epsilon))^k| \\
&\quad\leq 4k f_{\max}^k\epsilon + \frac{2f_{\max}}{\sqrt{N}} + \inf_\Gamma \mathbb{E}\,|(\mu_t^N(f_\epsilon))^k - (\mu_t^{(i),N}(f_\epsilon))^k|
\end{aligned}$$

Now by 7.4, since f_ϵ is compactly supported, for all N, there exists an $I(N, \epsilon)$ so that if $i \geq I(N, \epsilon)$ then

$$\inf_\Gamma \mathbb{E}\,|(\mu_t^N(f_\epsilon))^k - (\mu_t^{(i),N}(f_\epsilon))^k| \leq \frac{6f_{\max} + 1}{N},$$

giving that for $i \geq I(N, \epsilon)$

$$|\,\mathbb{E}(\mu_t(f))^k - \mathbb{E}(\mu_t^{(i)}(f))^k| \leq 4k f_{\max}^k\epsilon + \frac{2f_{\max}}{\sqrt{N}} + \frac{6f_{\max} + 1}{N}.$$

Choosing ϵ small enough and N big enough, this can be made arbitrarily small for $i \geq I(N, \epsilon)$, completing the proof. □

All that's left to show in proving 7.2 is our proof of 7.4, which is rather technical. The proof begins by constructing a coupling between μ_t^N and $\mu_t^{(i),N}$.

7.3 The Coupling

To couple the token process for two initial measures μ and ν, we make use of Strassen's classical representation of the Prokhorov distance, which itself metrizes weak convergence in $P(S)$. That is, the Prokhorov metric d_P on $P(S)$ can be interpreted for two measures μ and ν as

$$d_P(\mu,\nu) = \inf\{\epsilon\colon \exists \lambda \in \Lambda \text{ s.t. } \mathbb{P}_\lambda(d(\xi,\tilde{\xi}) > \epsilon) \leq \epsilon\},$$

where Λ is the set of all couplings $\lambda = (\xi, \tilde{\xi})$ with marginals μ and ν [18].

Fix $t > 0, N \geq 1$ and consider two initial measures μ and ν with $d_P(\nu,\mu) \leq \epsilon$ for some ϵ small. Write $\xi_i, 1 \leq i \leq N$ and $t_{ij}, 1 \leq i,j \leq N$ for the initial locations and meeting times of the token process generated from μ. Similarly write $\tilde{\xi}_i, 1 \leq i \leq N$ and $\tilde{t}_{ij}, 1 \leq i,j \leq N$ for the token process generated from ν. We will define a coupling of these to prove 7.4.

Intuitively, our coupling pairs the paths of each of N tokens between the two measures. By assumption, we can couple each token's starting location so ξ_i and $\tilde{\xi}_i$ are close with high probability. Then, whenever possible we couple all the meeting times, thus coupling the partition process of the tokens over time. Unfortunately, the inclusion of atomic measures - i.e. meetings at time $t = 0$ - means that a bit of care is needed to carry this out. Heuristically, we finesse this issue by smoothing out meetings at $t = 0$ by coupling them with meetings in some small interval $(0, t_*]$.

We will need the following standard lemma on exponential random variables.

Lemma 7.7. *Let $X \sim \exp(a)$ and $Y \sim \exp(b)$ be exponential random variables with rates a and b. Then there is a coupling of X and Y so that*

$$\mathbb{P}(X \neq Y) \leq 1 - \frac{\min(a,b)}{\max(a,b)}.$$

Proof. This follows from the simple calculation of integrating the overlap in the two densities, i.e.

$$\int_0^\infty \min(a\exp(-at), b\exp(-bt))dt = \frac{\min(a,b)}{\max(a,b)}.$$

□

Our coupling between the token processes μ_t^N and ν_t^N is then defined as follows. First, by definition of the Prokhorov distance we may couple ξ_i and $\tilde{\xi}_i$ for each $1 \leq i \leq N$ such that

$$\mathbb{P}(d(\xi_i, \tilde{\xi}_i) > \epsilon) \leq \epsilon$$

where $\epsilon \geq d_P(\mu,\nu)$.

Then, for every pair i,j with both $d(\xi_i,\xi_j) > 0$ and $d(\tilde{\xi}_i,\tilde{\xi}_j) > 0$, we couple the meeting times t_{ij} and $\tilde{t}_{ij}$ as in Lemma 7.7. In the case $d(\xi_i,\xi_j) = 0$ we have $t_{ij} = 0$ and there is no possibility of a better than independent coupling with (a non-zero) $\tilde{t}_{ij}$. That is, in this case all couplings of t_{ij} with $\tilde{t}_{ij}$ are equivalent. We similarly handle the case $d(\tilde{\xi}_i,\tilde{\xi}_j) = 0$.

7.4 Proof of 7.4

We being by recalling the (random) equivalence relation $\sim_0$ on $1, 2, \ldots, N$ given by $i \sim_0 j$ if $\xi_i = \xi_j$. For simplicity of notation we'll shorten $\sim_0$ to $\sim$.

We say that an outcome of the coupling is $(\epsilon, d_1, d_2, t_*)$-**good** for $0 < d_1 < d_2 < \infty$ if

(G1) For all i, j either $i \sim j$ or $d(\xi_i, \xi_j) \geq d_1$,

(G2) For all i, $d(\xi_i, \tilde{\xi}_i) \leq \epsilon$,

(G3) For all i, j s.t. $i \nsim j$ we have $t_{ij} = \tilde{t}_{ij}$,

(G4) For all i, j s.t. $i \nsim j$ we have $t_{ij} \geq t_* > 0$,

(G5) For all i, j s.t. $i \sim j$ we have $\tilde{t}_{ij} \leq t_*$,

(G6) For all i, j we have $d(\xi_i, \xi_j) \leq d_2$ and $d(\tilde{\xi}_i, \tilde{\xi}_j) \leq d_2$.

In words, we allow the possibility that some tokens i and j meet at $t_{ij} = 0$ even if $\tilde{t}_{ij} \neq 0$. However, we insist that all such meetings $\tilde{t}_{ij}$ happen before t_* and no other meetings happen before t_*. As part of this, we require tokens i and j with $\xi_i \neq \xi_j$ to be spaced by at least d_1, and require that no pair of tokens for either process are further than d_2. Then, our outcome is **good** if in addition, all the meetings happening after t_* are coupled.

Our motivation for this definition is that in the case that the outcome is $(\epsilon, d_1, d_2, t_*)$-**good** it becomes simple to bound

$$|\mu_t^N(f) - \nu_t^N(f)|,$$

for compactly supported f. Consider a fixed, compactly supported $f \in C_b(S)$. Then by compactness f must be uniformly continuous, so for every $\alpha > 0$ there is a $\delta_f(\alpha) > 0$ so that for all $x, y \in S$, if $d(x, y) \leq \delta_f(\alpha)$ then

$$|f(x) - f(y)| \leq \alpha.$$

Lemma 7.8. *If an outcome of the coupling is $(\delta_f(\alpha), d_1, d_2, t_*)$-**good**, then for $t \geq t_*$*

$$|\mu_t^N(f) - \nu_t^N(f)| \leq \alpha.$$

Proof. First we examine the process at time t_*. By $(G2)$, for μ^N all meetings t_{ij} happen either at time $t = 0$ or at times $t \geq t_*$. By $(G3), (G4)$ and $(G5)$ the only meetings of ν^N that occur before t_* are those between tokens i and j with $t_{ij} = 0$, and all such meetings occur before t_*.

From this we will be able to show that by time t_*, the two partitions processes of tokens are the same. Let I_i be the equivalence class of a token i under $\sim_0$. By definition, then for μ^N all tokens in I_i meet at time $t = 0$ and have no other meetings until at least time t_* and so

$$u_i(t_*) = \inf I_i,$$

since all tokens in I_i are absorbed into token $\inf I_i$ instantly at $t = 0$. Recall here that $u_i(t)$ is the owner of token i at time t. Note this trivially true if i is a singleton under $\sim$.

By assumption, for ν^N all pairs of tokens in I_i meet before t_* and have no meetings with tokens outside of I_i until after t_*. These meetings could come in any order but by time t_* all tokens of ν^N in a class I_i are also owned by $\inf I_i$. That is, we can't say exactly what meeting tree occurs, but we know the final result.

Therefore, for all $1 \leq i \leq N$ we have

$$u_i(t_*) = \tilde{u}_i(t_*). \tag{9}$$

We claim that Equation (9) also holds for all times $t \geq t_*$. This however follows easily from $(G3)$ as the only meetings for μ^N that occur after t_* are coupled to the corresponding meetings for ν^N. Thus once the two token partitions are synced at t_*, they remain so for all times after.

Therefore, applying $(G2)$ we then have that for all tokens $1 \leq i \leq N$ and times $t \geq t_*$

$$d(\xi_{u_i(t)}, \tilde{\xi}_{\tilde{u}_i(t)}) \leq \delta(\alpha).$$

This then implies that

$$|f(\xi_{u_i(t)}) - f(\tilde{\xi}_{\tilde{u}_i(t)})| \leq \alpha,$$

for $t \geq t_*$.

Recalling Section 2.2 we can conclude that for all times $t \geq t_*$

$$|\mu_t^N(f) - \nu_t^N(f)| \leq \sum_{i=1}^{N} \frac{1}{N} |f(\xi_{u_i(t)}) - f(\tilde{\xi}_{\tilde{u}_i(t)})| \leq \alpha.$$

□

The rest of the difficulty in our proof of 7.4 is calculating the probability that the outcome is $(\epsilon, d_1, d_2, t_*)$-**good**, or more specifically bounding the probability that it is not. To start with, we need a technical lemma on the smoothness of the rate function $\phi(x)$. This can be more easily avoided with some smoothness conditions on ϕ but we prefer to give a proof in more generality.

Lemma 7.9. *Fix $0 < d_1 \leq d_2 < \infty$. There exists a non-decreasing $G_{d_1,d_2} \colon \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ such that*

$$1 - \frac{\min(\phi(x), \phi(y))}{\max(\phi(x), \phi(y))} \leq G_{d_1,d_2}(|x - y|)$$

for all $x, y \in [d_1, d_2]$ and also $\lim_{z \downarrow 0} G_{d_1,d_2}(z) = 0$.

Proof. Write $C = [d_1, d_2]^2$ and define

$$G_{d_1,d_2}(z) = \sup_{(x,y) \in C \colon |x-y| \leq z} 1 - \frac{\min(\phi(x), \phi(y))}{\max(\phi(x), \phi(y))}.$$

Clearly $G_{d_1,d_2}(z)$, as the supremum on a compact set of a function bounded by 1, is well defined and itself satisfies $0 \leq G_{d_1,d_2}(z) \leq 1$ for all z. Also clear is that $G_{d_1,d_2}(0) = 0$ and $G_{d_1,d_2}(z)$ is non-decreasing in z, since as z increases so does the domain of the supremum. So we need only show that $\lim_{z \downarrow 0} G_{d_1,d_2}(z) = 0$.

Assume otherwise and so there exists an $\epsilon > 0$ and a sequence $z_i \downarrow 0$ such that $G_{d_1,d_2}(z_i) \geq \epsilon$. By the compactness of $\{(x, y) \in C \colon |x - y| \leq z\}$ there exists

a sequence of points $(x_i, y_i) \in C$ with $|x_i - y_i| \leq z_i$ attaining the supremum, that is

$$1 - \frac{\min(\phi(x_i), \phi(y_i))}{\max(\phi(x_i), \phi(y_i))} = G_{d_1,d_2}(z_i) \geq \epsilon.$$

By compactness again, there exists a limiting subsequence (x_j, y_j) converging to some point $(x_\infty, y_\infty) \in C$. As $z_j \downarrow 0$, we must have $|x_j - y_j| \downarrow 0$ and so $|x_\infty - y_\infty| = 0$, i.e. $x_\infty = y_\infty$. Since

$$1 - \frac{\min(\phi(x), \phi(y))}{\max(\phi(x), \phi(y))}$$

is continuous for (x, y) in C - as ϕ is bounded away from zero for $x, y \in [d_1, d_2]$ - this implies that

$$1 - \frac{\min(\phi(x_\infty), \phi(y_\infty))}{\max(\phi(x_\infty), \phi(y_\infty))} = \lim_j G_{d_1,d_2}(z_j) \geq \epsilon,$$

clearly a contradiction as it must be 0. □

We'll need to introduce a bit of temporary notation before moving forward. Write

$$\phi_{\max}(a, b) = \sup_{x \in [a,b]} \phi(x)$$

which by the continuity of ϕ is finite for all $0 < a \leq b < \infty$.

Also write

$$\phi_{\min}(z) = \inf_{x \in (0,z]} \phi(x)$$

for $z < \infty$, which by the assumption $\lim_{x \downarrow 0} \phi(x) = \infty$ must be non-zero. Note the difference in the two domains being optimized over.

Also, for $z > 0$, we write

$$F_\mu(z) = \mathbb{P}(0 < d(\xi_1, \xi_2) \leq z),$$

which by right-continuity must satisfy $F_\mu(z) \downarrow 0$ as $z \downarrow 0$. Also write

$$\bar{F}_\mu(z) = \mathbb{P}(z \leq d(\xi_1, \xi_2)),$$

which satisfies $\bar{F}_\mu(z) \downarrow 0$ as $z \uparrow \infty$. Also clearly

$$F_\mu(z) + \bar{F}_\mu(z) + \mathbb{P}(d(\xi_1, \xi_2) \in \{0, z\}) = 1.$$

We're finally to state our initial bound on the probability that the coupling is good.

Lemma 7.10. *Consider initial measures μ and ν with $d_P(\mu, \nu) \leq \frac{\epsilon^2}{2}$. If $\epsilon \leq \frac{d_1}{5}$ then the probability that the outcome is not $(\epsilon, d_1, d_2 t_*)$-**good** is bounded as follows:*

$$\begin{aligned}\mathbb{P}(\textit{Not } (\epsilon, d_1, d_2, t_*) - \textit{good}) \leq & \binom{N}{2} F_\mu(d_1) + N\epsilon + \binom{N}{2} G_{\frac{d_1}{2},d_2}(2\epsilon) \\ & + \binom{N}{2}(1 - \exp(-\phi_{\max}(d_1, d_2)t_*)) + \binom{N}{2} \exp(-\phi_{\min}(2\epsilon)t_*) \\ & + \binom{N}{2} \bar{F}_\mu(\frac{d_2}{2}).\end{aligned}$$

To prove Lemma 7.10, we first bound the probability of failure for each of the requirements for the coupling to be $(\epsilon, d_1, d_2, t_*)$-good. For each of these lemmas, assume that μ and ν satisfy the conditions of Lemma 7.10.

Lemma 7.11.

$$\mathbb{P}\left(Not\ (G1)\right) \leq \binom{N}{2} F_\mu(d_1).$$

Proof. For any pair i, j, we have that the probability $(G1)$ fails is

$$\mathbb{P}(i \nsim j \text{ and } d(\xi_i, \xi_j) < d_1) \leq \mathbb{P}(d(\xi_i, \xi_j) \in (0, d_1]) = F_\mu(d_1).$$

The lemma then follows by a simple union bound. □

Lemma 7.12.

$$\mathbb{P}\left(Not\ (G2)\right) \leq N\epsilon.$$

Proof. For each i, by the definition of the coupling we have

$$\mathbb{P}(d(\xi_i, \tilde{\xi}_i) > \epsilon) \leq \epsilon.$$

Therefore by a simple union bound

$$\mathbb{P}(\exists i \text{ s.t. } (G1) \text{ fails for } i) \leq N\epsilon.$$

□

Write $\sigma(\Xi)$ for the sigma field of the initial locations of the tokens, that is

$$\sigma(\Xi) = \sigma(\xi_i, \tilde{\xi}_i, 1 \leq i \leq N).$$

Lemma 7.13.

$$\mathbb{P}\left((G1), (G2), (G6), Not\ (G3)\right) \leq \binom{N}{2} G_{\frac{d_1}{2}, d_2}(2\epsilon).$$

Proof. For simplicity of notation, write

$$d_{ij} = d(\xi_i, \xi_j), \qquad \tilde{d}_{ij} = d(\tilde{\xi}_i, \tilde{\xi}_j).$$

First when $(G2)$ holds, by the triangle inequality we always have

$$|d_{ij} - \tilde{d}_{ij}| \leq d(\xi_i, \tilde{\xi}_i) + d(\xi_j, \tilde{\xi}_j) \leq 2\epsilon.$$

When $(G1)$ holds, if $i \nsim j$ then

$$d_{ij} \geq d_1$$

and so as $\epsilon \leq \frac{d_1}{5}$

$$\tilde{d}_{ij} \geq d_1 - 2\epsilon \geq \frac{d_1}{2},$$

giving both d_{ij} and $\tilde{d}_{ij}$ at least $\frac{d_1}{2}$.

By Lemma 7.7, for each such pair i, j the probability that $(G3)$ fails given give the initial locations satisfies

$$\mathbb{P}(t_{ij} \neq \tilde{t}_{ij}|\Xi) \leq 1 - \frac{\min(\phi(d_{ij}), \phi(\tilde{d}_{ij}))}{\max(\phi(d_{ij}), \phi(\tilde{d}_{ij}))}.$$

When $(G6)$ also holds, we have

$$\frac{d_1}{2} \leq d_{ij}, \tilde{d}_{ij} \leq d_2.$$

Applying Lemma 7.9 with $\frac{d_1}{2}$ and d_2, using that $(G1), (G2)$ and $(G6)$ are determined by the initial locations and so in $\sigma(\Xi)$, we have that

$$\begin{aligned}
\mathbb{P}\left(t_{ij} \neq \tilde{t}_{ij}, (G1), (G2), (G6)\right) &= \mathbb{E}\,\mathbb{P}\left(t_{ij} \neq \tilde{t}_{ij}, (G1), (G2), (G6)|\Xi\right) \\
&= \mathbb{E}\left(1\left((G1), (G2), (G6)\right)\mathbb{P}(t_{ij} \neq \tilde{t}_{ij}|\Xi)\right) \\
&\leq \mathbb{E}\left(1\left((G1), (G2), (G6)\right)\left(1 - \frac{\min(\phi(d_{ij}), \phi(\tilde{d}_{ij}))}{\max(\phi(d_{ij}), \phi(\tilde{d}_{ij}))}.\right)\right) \\
&\leq \mathbb{E}\left(1\left((G1), (G2), (G6)\right)G_{\frac{d_1}{2}, d_2}(|d_{ij} - \tilde{d}_{ij}|)\right) \\
&\leq \mathbb{E}\left(1\left((G1), (G2), (G6)\right)G_{\frac{d_1}{2}, d_2}(2\epsilon)\right) \\
&\leq G_{\frac{d_1}{2}, d_2}(2\epsilon)
\end{aligned}$$

recalling that $G_{\frac{d_1}{2}, d_2}$ is non-decreasing and $|d_{ij} - \tilde{d}_{ij}| \leq 2\epsilon$.

Thus we have that

$$\mathbb{P}\left((G1), (G2), (G6) \text{ but } (G3) \text{ fails for } i, j\right) \leq G_{\frac{d_1}{2}, d_2}(2\epsilon)$$

and the proof follows by a simple union bound. □

Lemma 7.14.

$$\mathbb{P}\left((G1), (G6), \textit{Not } (G4)\right) \leq \binom{N}{2}(1 - \exp(-\phi_{\max}(d_1, d_2)t_*)).$$

Proof. When $(G1)$ and $(G6)$ hold, if $i \nsim j$ then $d_1 \leq d(\xi_i, \xi_j) \leq d_2$. Therefore by definition of $\phi_{\max}(d_1, d_2)$ we have

$$\phi(d(\xi_i, \xi_j)) \leq \phi_{\max}(d_1, d_2).$$

So the probability that $(G4)$ fails for i, j can be calculated, using that $(G1)$ and $(G6)$ are in $\sigma(\Xi)$, and that t_{ij} is (conditionally) exponentially distributed, by

$$\begin{aligned}
\mathbb{P}(t_{ij} < t_*, (G1), (G6)) &= \mathbb{E}\left(1\left((G1), (G6)\right)\mathbb{P}(t_{ij} < t_*|\Xi)\right) \\
&= \mathbb{E}\left(1\left((G1), (G6)\right)(1 - \exp(-\phi(d(\xi_i, \xi_j))t_*))\right) \\
&\leq \mathbb{E}\left(1\left((G1), (G6)\right)(1 - \exp(-\phi_{\max}(d_1, d_2)t_*))\right) \\
&\leq 1 - \exp(-\phi_{\max}(d_1, d_2)t_*).
\end{aligned}$$

A simple union bound completes the proof. □

Lemma 7.15.

$$\mathbb{P}\left((G1),(G2),\mathit{Not}\ (G5)\right) \leq \binom{N}{2}\exp(-\phi_{\min}(2\epsilon)t_*).$$

Proof. When $(G1)$ and $(G2)$ hold, if $i \sim j$, using $\epsilon \leq \frac{d_1}{5}$, from the triangle inequality we have

$$d(\tilde{\xi}_i, \tilde{\xi}_j) \leq 2\epsilon.$$

Therefore

$$\phi(d(\tilde{\xi}_i, \tilde{\xi}_j)) \geq \phi_{\min}(2\epsilon),$$

and so the probability that $(G5)$ fails for i, j, using that $(G1)$ and $(G2)$ are in $\sigma(\Xi)$ is

$$\begin{aligned}\mathbb{P}\left(\tilde{t}_{ij} \geq t_*, (G1), (G2)\right) &= \mathbb{E}\left(1\left((G1),(G2)\right)\mathbb{P}(\tilde{t}_{ij} \geq t_* | \Xi)\right) \\ &\leq \mathbb{E}\left(1\left((G1),(G2)\right)\exp(-\phi(d(\tilde{\xi}_i, \tilde{\xi}_j))t_*)\right) \\ &\leq \exp(-\phi_{\min}(2\epsilon)t_*).\end{aligned}$$

The lemma then follows from a simple union bound. □

Lemma 7.16.

$$\mathbb{P}\left((G2),\mathit{Not}\ (G6)\right) \leq \frac{N}{2}\bar{F}_\mu(\frac{d_2}{2}).$$

Proof. First, by definition we have that

$$\mathbb{P}(d(\xi_1, \xi_2) \geq \frac{d_2}{2}) = \bar{F}_\mu(\frac{d_2}{2}).$$

Therefore the probability that

$$d(\xi_i, \xi_j) \leq \frac{d_2}{2} \tag{10}$$

fails for any pair ξ_i, ξ_j can be bounded by the union bound

$$\frac{N}{2}\bar{F}_\mu(\frac{d_2}{2}).$$

Next, by assumption

$$d_2 \geq d_1 > 4\epsilon$$

and so if Equation (10) holds for all i and j, then by the triangle inequality

$$\begin{aligned}d(\tilde{\xi}_i, \tilde{\xi}_j) &\leq d(\xi_i, \tilde{\xi}_i) + d(\xi_j, \tilde{\xi}_j) + d(\xi_i, \xi_j) \\ &\leq 2\epsilon + \frac{d_2}{2} \\ &\leq d_2\end{aligned}$$

showing that $(G2)$ holds for all $\tilde{\xi}_i$ and $\tilde{\xi}_j$ completing the proof. □

We can now easily complete the proof of Lemma 7.10.

Proof. To bound the probability that the outcome of the coupling is not $(\epsilon, d_1, d_2 t_*)$-**good**, by a union bound we can examine the probability that each of the conditions fails. Since

$$\begin{aligned}\{\text{Not Good}\} \subset &\{\text{Not } (G1)\} \cup \{\text{Not } (G2)\} \cup \{(G1), (G2), \text{Not } (G5)\} \\ &\cup \{(G2), \text{Not } (G6)\} \cup \{(G1), (G2), (G6), \text{Not } (G3)\} \\ &\cup \{(G1), (G6), \text{Not } (G4)\}\end{aligned}$$

we can combine Lemmas 7.11 to 7.16 to complete the proof. □

Note that Lemma 7.10 is nowhere near an optimal bound on the probability the coupling is not good. Our approach here follows the philosophy that since the probability that the coupling isn't $(\epsilon, d_1, d_2, t_*)$-good approaches zero, even a "worst case" bound on it will suffice, which is exactly what happens.

We next need to show that the probability that our outcome is not $(\epsilon, d_1, d_2, t_*)$-good can be made arbitrarily small if we assume $d_P(\mu, \nu)$ is small enough.

Lemma 7.17. *There exists a choice of d_1, d_2, t_* and $\epsilon_1(N, \phi, \mu) \geq 0$- depending on N, ϕ, μ and (S, d) - such that for any $\epsilon \leq \epsilon_1(N, \phi, \mu)$ and all measures $\nu \in P(S)$, if $d_P(\mu, \nu) \leq \frac{\epsilon^2}{2}$ then*

$$\mathbb{P}(\textit{Not } (\epsilon, d_*, t_*) - \textit{good}) \leq \frac{6}{N}.$$

Proof. We will prove this by showing that there exists a choice of constants d_1, d_2, t_*, and ϵ - picked in that order - which applied to Lemma 7.10 satisfy our bound.

First, as $\lim_{z \downarrow 0} F_\mu(z) = 0$ we may pick $d_1 > 0$ so that

$$F_\mu(d_1) \leq \frac{1}{N} \binom{N}{2}^{-1},$$

and thus

$$\binom{N}{2} F_\mu(d_1) \leq \frac{1}{N}. \tag{11}$$

Next, as $\bar{F}_\mu(z) \downarrow 0$ as $z \uparrow \infty$, we can pick d_2 large enough so that

$$\binom{N}{2} \bar{F}_\mu(d_2) \leq \frac{1}{N}.$$

For t_*, as $\phi_{\max}(d_1, d_2) < \infty$, we may pick $t_* > 0$ so that

$$t_* \leq -\frac{\ln\left(1 - \frac{1}{N\binom{N}{2}}\right)}{\phi_{\max}(d_1, d_2)},$$

giving that

$$\binom{N}{2} (1 - \exp(-\phi_{\max}(d_1, d_2) t_*)) \leq \frac{1}{N}. \tag{12}$$

Next, as $\lim_{x\downarrow 0}\phi(x)=\infty$, there exists a γ such that for all $0<x\leq\gamma$ we have

$$\phi(x)\geq\frac{\ln(N\binom{N}{2})}{t_*}.$$

Similarly, as $\lim_{z\downarrow 0}G_{\frac{d_1}{2},d_2}(z)=0$, there exists a β so that if $x<\beta$ then

$$G_{\frac{d_1}{2},d_2}(2x)\leq N^{-1}\binom{N}{2}^{-1}.$$

Therefore, we take

$$\epsilon_1(N,\phi,\mu)=\min\left(\frac{1}{N^2},\frac{d_1}{5},\beta,\frac{\gamma}{2}\right).$$

This will allow us to bound all the terms in Lemma 7.10. Note that this also easily implies - as $d_1\leq d_2$ - that $\epsilon_1\leq\frac{d_2}{4}$, satisfying the minor condition of Lemma 7.10. First, if $\epsilon\leq\epsilon_1(N,\phi,\mu)$ we easily have

$$\epsilon\leq\frac{d_1}{5},$$

as well as

$$N\epsilon\leq\frac{1}{N}. \tag{13}$$

Using $\epsilon\leq\beta$ we have that

$$G_{\frac{d_1}{2},d_2}(2\epsilon)\leq N^{-1}\binom{N}{2}^{-1}$$

and so

$$\binom{N}{2}G_{\frac{d_1}{2},d_2}(2\epsilon)\leq\frac{1}{N}. \tag{14}$$

Finally as $2\epsilon\leq\gamma$ we have

$$\phi(2\epsilon)\geq\frac{\ln(N\binom{N}{2})}{t_*}$$

and so

$$\binom{N}{2}\exp(-\phi(2\epsilon)t_*)\leq\frac{1}{N}. \tag{15}$$

Applying Equations (11) to (15) to Lemma 7.10, we see that if $d(\mu,\nu)\leq\frac{\epsilon^2}{2}$ for some $\epsilon\leq\epsilon_1$ then

$$\mathbb{P}(\text{Not }(\epsilon,d_1,d_2,t_*)-\text{good})\leq\frac{6}{N}.$$

□

We are now able finish our proof of 7.4.

Proof. By Lemma 7.17, there exists a choice of d_1, d_2, t_* and $\epsilon_1(N, \phi, \mu)$ such that for any $\epsilon \leq \epsilon_1(N, \phi, \mu)$, if

$$d_P(\mu, \nu) \leq \frac{\epsilon^2}{2}$$

then

$$\mathbb{P}(\text{Not } (\epsilon, d_1, d_2, t_*) - \text{good}) \leq \frac{6}{N}.$$

Recall $\delta_f(x)$ as in Lemma 7.8 and define ϵ_0 by

$$\epsilon_0 = \min(\epsilon_1(N, \phi, \mu), \delta_f(\frac{1}{N})).$$

Therefore, calculating the expectation of

$$|\mu_t^N(f) - \nu_t^N(f)|$$

whenever $d_P(\mu, \nu) \leq \epsilon_0$, using our coupling and Lemma 7.8 with $\epsilon_0 \leq \delta_f(\frac{1}{N})$ we have

$$\begin{aligned}
\mathbb{E}\,|\mu_t^N(f) - \nu_t^N(f)| &\leq \mathbb{E}\,|\mu_t^N(f) - \nu_t^N(f)|\,(1(\text{good}) + 1(\text{Not good}) \\
&\leq \frac{1}{N}\,\mathbb{P}(\text{good}) + \mathbb{E}\,f_{\max}\,\mathbb{P}(\text{Not } (\epsilon, d_1, d_2, t_*) - \text{good}) \\
&\leq \frac{1}{N} + f_{\max}\frac{6}{N} \\
&\leq \frac{6f_{\max} + 1}{N}
\end{aligned}$$

where $f_{\max}$ is the maximum of $|f|$ on S.

This completes the proof, since by assumption $\mu^{(i)} \to \mu$ weakly and so for i large enough

$$d_P(\mu^{(i)}, \mu) \leq \epsilon_0.$$

□

7.5 Moment Calculations for $\mu_t(f)$

In this section we mention briefly a method that gives an explicit closed form calculation for the moments of $\mu_t(f)$ for any time $t \geq 0$, fixed $f \in C_b(S)$, and initial measure $\mu \in P(S)$. In Section 6.2 we calculate the second moment of $\mu_t(f)$ by conditioning on μ_0^N, summing over the different possible meeting trees on 2 tokens giving the expectation of $\left(\mu_t^N\right)^2$, and then taking the $N \to \infty$ limit.

In theory, this same approach gives a perhaps simpler proof of 7.2 as compared to the coupling method of Section 7.3, by showing explicitly that if $\mu^{(i)} \to \mu$ weakly in $P(S)$ then

$$\mathbb{E}\left(\mu_t^{(i)}(f)\right)^k \to \mathbb{E}\left(\mu_t(f)\right)^k,$$

for all $k \geq 1$. The calculation for $k > 2$ follows the same procedure as for the second moment: first conditioning on μ_0^N, then summing over possible meeting trees on k tokens, and finally taking the $N \to \infty$ limit. In fact, tighter control of

the moments $\mu_t^N(f)$ of the N token processes could even provide an alternative proof of the existence of the limit MC process $\mu_t, t \geq 0$ (as in 2.1) without resort to Kingman's Paintbox theorem.

In practice however, for anything past the second moment this formula appears to be intractable, necessitating our proof of 7.2 by coupling methods. The main technical issue is two-fold. First, the number of possible meeting trees on k tokens increases approximately as $\binom{k}{2}!$ making an explicit computation (in the form of 6.5) impossible. More importantly, the (necessary) assumption that

$$\lim_{x \downarrow 0} \phi(x) = \infty$$

means that for any fixed meeting tree T, the probability that at time t the observed meeting tree $S(t)$ is $S(t) = T$ does not depend continuously (in the weak topology) on the initial measure μ. A posteriori of course, by the coupling argument of Section 7.3, we know that the moments of $\mu_t(f)$ depend continuously on the initial measure μ nonetheless. We leave as an open question the details of the pursuit of an alternative approach to Item 2 by moment methods.

8 Continuity and Uniqueness

In this section we prove first that from any initial measure, the MC process is cadlag almost surely. This builds on 6.1 where we've shown that the MC process is almost surely cadlag at $t = 0$. Here we're interested in the global path properties of $\mu_t, t \geq 0$.

Proposition 8.1. *From any initial measure μ, the MC process $\mu_t, t \geq 0$ is cadlag in $P(S)$ with respect to the weak topology, almost surely.*

From this, we are then able to complete our proof of Item 2, by proving that our extension of the MC process is unique.

8.1 Proof of 8.1

Our proof of 8.1 begins by considering the Total Variation distance between the measure valued processes μ_t^N and μ_t^M for some positive time t. Write d_{TV} for the Total Variation metric on $P(S)$.

Lemma 8.2. *Fix $t_0 > 0$ and $M \geq N$. Then*

$$\sup_{t \in [t_0, \infty)} d_{\mathrm{TV}}(\mu_t^N, \mu_t^M) \leq d_{\mathrm{TV}}(\mu_{t_0}^N, \mu_{t_0}^M).$$

Proof. The proof follows from the simple idea that the total variation distance can only decrease after each merger (i.e. jump) of the process $\mu_t^M, t \geq t_0$. To see this, write

$$\mu_{t_0}^M = \sum_{k=1}^{K^M(t_0)} q_k \delta(s_k)$$

with the support written in any order. Clearly $\operatorname{supp}\mu_{t_0}^N \subset \operatorname{supp}\mu_{t_0}^M$, since its given by the empirical distribution of the tokens, so we can also write

$$\mu_{t_0}^N = \sum_{k=1}^{K^M(t_0)} \tilde{q}_k \delta(s_k)$$

for some masses $\tilde{q}_k \geq 0$. Note that we assume no particular relation between q_k and $\tilde{q}_k$ - the point is that they are co-located.

The total variation distance of $\mu_{t_0}^N$ and $\mu_{t_0}^M$ is then given by

$$d_{\mathrm{TV}}(\mu_{t_)}^N, \mu_{t_0}^M) = \frac{1}{2} \sum_{k=1}^{K^M(t_0)} |q_k - \tilde{q}_k|.$$

Consider the first jump time T of μ^M. Clearly μ_t^N can't jump on $[t_0, T)$ if μ_t^M doesn't. For simplicity of notation assume without loss of generality that at time T the mass at s_1 merges into s_2. Then

$$\begin{aligned} 2d_{\mathrm{TV}}(\mu_T^N, \mu_T^M) &= |(q_1 + q_2) - (\tilde{q}_1 + \tilde{q}_2)| + \sum_{k=3}^{K^M(t_0)} |q_k - \tilde{q}_k| \\ &\leq |q_1 - \tilde{q}_1| + |q_2 + \tilde{q}_2| + \sum_{k=3}^{K^M(t_0)} |q_k - \tilde{q}_k| \\ &= 2d_{\mathrm{TV}}(\mu_{t_0}^N, \mu_{t_0}^M). \end{aligned}$$

As this same argument holds at each of the (finitely many) subsequent jumps, inducting over the meeting times of $\mu_t^M, t \geq t_0$ the proof is complete. □

Note that the above proof of Lemma 8.2 holds at $t = 0$, however as there is no longer necessarily a countable support representation of $\mu_0 = \mu$, we don't know if $\mu_0^N, N \geq 1$ is Cauchy (w.r.t. d_{TV}) in $P(S)$. As in Section 2.6, we know $\mu_0^N \to \mu_0$ weakly, but for non-atomic μ_0 not in Total Variation distance.

Now, writing $D_I(P(S))$ for the space of cadlag paths from an interval I to $P(S)$, viewed under the uniform (total variation) norm. That is, for any two functions $f, g \in D_I(P(S))$, we write their uniform distance as

$$d_{\mathrm{unif}}(f, g) = \sup_{x \in I} d_{\mathrm{TV}}(f(x), g(x)).$$

Note that typically for a metric space E, the space $D_I(E)$ is more naturally considered with respect to the Skorohod metric, which allows paths to "wiggle" in both space and time. However, as μ_t^N and μ_t are intrinsically coupled, so are their jump times and so there is no need for the flexibility in time provided by the Skorohod metric. This allows us to use the much stronger topology of uniform convergence, which among other advantages is complete (see Section 3.12 [7]).

Lemma 8.3. *For any complete metric space E and any interval I, the space $D_I(E)$ is complete under the uniform norm.*

As the empirical measures processes $\mu_t^N, t \geq 0$ are finite jump processes, they are clearly in $D_I(P(S))$ for any interval I. Recalling Lemma 2.12, Lemma 8.2 implies that for any fixed $t_0 > 0$ the sequence $\mu_t^N, t \geq t_0$ is Cauchy in $D_{[t_0,\infty)}(P(S))$ under the uniform norm. Lemma 8.3 then tells us that the sequence has a cadlag limit.

As we already know (2.1) that $\mu_t, t \geq t_0$ is the pointwise limit of the sequence $\mu_t^N, t \geq t_0$, the following lemma is immediate.

Lemma 8.4. *For any fixed $t_0 > 0$, $\mu_t, t \geq t_0$ is cadlag almost surely.*

We can now finish our proof of 8.1.

Proof. Let A_n be the event that $\mu_t, t \geq \frac{1}{N}$ is cadlag. By Lemma 8.4, we have that

$$\mathbb{P}(A_n) = 1.$$

It then follows that

$$\mathbb{P}(\cap_n A_n) = 1,$$

and so almost surely μ_t is cadlag on all of $(0, \infty)$. Since the topology of the Total Variation distance is stronger than the topology of weak convergence, this immediately implies that μ_t is also cadlag with respect to the weak topology on $P(S)$.

By 6.1, we know that μ_t is also right continuous at 0 in the weak topology almost surely, completing the proof. □

8.2 Uniqueness of the Extension

We are now ready to complete the proof of Item 2. Writing $D_{\mathbb{R}_{\geq 0}}(P(S))$ for the space of cadlag maps from $\mathbb{R}_{\geq 0}$ to $P(S)$. By 8.1, for each initial measure μ we can think of the MC process as being given by a measure $\mathbb{P}_\mu$ on $D_{\mathbb{R}}(P(S))$. Our constructed process is then completely determined by this family of measures

$$\mathbb{P}_{\text{MC}} = \{\mathbb{P}_\mu, \mu \in P(S)\}.$$

To complete our proof of Item 2 we need only show that $\mathbb{P}_{\text{MC}}$ is the unique such family of measures on $D_{\mathbb{R}}(P(S))$.

Proposition 8.5. *Any other cadlag, Feller continuous extension of the Metric Coalescent satisfying the properties of Item 2 is identically distributed to $\mathbb{P}_{\text{MC}}$.*

Proof. This follows easily from a standard argument given that our process is cadlag, Markov and Feller continuous. The key idea is that any cadlag time-homogeneous Markov process is separable, i.e. determined uniquely by its finite dimensional distributions (fdd's). The Chapman-Kolmogorov equation show that the fdd's are determined by the single dimensional distributions

$$P_t(x, A) = \mathbb{P}(\mu_t \in A | \mu_0 = x),$$

for $A \subset P(P(S))$. Feller continuity implies then that if $\mu^{(i)} \to \mu$ so do the single dimensional distributions and therefore so do all the fdd's of the processes.

Since any other extension of the MC agrees with $\mathbb{P}_{\text{MC}}$ on all of $P_{\text{fs}}(S)$ which is dense in $P(S)$ under the weak topology, we can conclude that any other such extension has the same fdd's and is therefore identically distributed on $D_{\mathbb{R}}(P(S))$. □

9 Some Motivating Examples

Having finished our proof of Item 2, in this section we look at a few particular examples of the Metric Coalescent that motivate some of the assumptions we have made. In particular, we show that without assumptions on the rate function $\phi(x)$, even in the world of finitely supported measures, the Metric Coalescent would not be Feller continuous with respect to the weak topology. We'll also give a construction showing that even for countably supported measures in $P(S)$, without the assumption of compact support, the MC process need not be in $P_{\text{fs}}(S)$ for positive times.

9.1 The Rate $\phi(x)$ Near Zero

Much of the technical complexity in the proof of Item 2 is a result of $\phi(x)$ being unbounded at $x = 0$. A simple example however shows this to be necessary for Feller continuity to hold. Intuitively, the issue can be seen by looking at a converging sequence $x_i \to x_\infty$ and measures

$$\mu^{(i)} = \frac{1}{2}\delta(x_i) + \frac{1}{2}\delta(x_\infty)$$

and $\mu^\infty = \delta(x_\infty)$. Clearly $\mu^{(i)} \to \mu^\infty$ weakly, however if $\phi(d(x_i, x_\infty))$ doesn't increase to ∞, then the time between the two atoms merging (for some subsequence) is bounded below.

To actually construct a working counter-example, we need to consider at least three atoms.

9.1.1 An Example

For ease of notation we work on $S = [0, 1]$ with the standard metric - however this is by no means necessary. We will construct a sequence of initial measures $\mu^{(n)} \to \mu^\infty$ and a continuous function $f\colon [0, 1] \to \mathbb{R}$ such that

$$\mu_t^{(n)}(f) \not\Rightarrow \mu_t^\infty(f),$$

for $t > 0$ whenever $\lim_{x\downarrow 0} \phi(x) \neq \infty$.

Let $x_n \downarrow 0$ be a sequence of points with the assumption - again purely for simplicity of notation - that $0 \leq x_n \leq \frac{1}{2}$ and consider the sequence of measures

$$\mu^{(n)} = \frac{1}{3}\left(\delta(0) + \delta(1) + \delta(x_n)\right).$$

Clearly the sequence μ^n limits weakly to

$$\mu^\infty = \frac{2}{3}\delta(0) + \frac{1}{3}\delta(1).$$

Let $f(x)$ be any continuous function supported near 1, for instance

$$f(x) = \begin{cases} 0 & \text{if } 0 \leq x \leq \frac{1}{2}, \\ 2x - 1 & \text{if } \frac{1}{2} \leq x \leq 1. \end{cases}$$

As in Section 6 we will consider the martingales $\mu_t^{(n)}(f)$ and similarly $\mu_t^\infty(f)$.

First consider $\mu_t^\infty(f)$. Let $T_{0,1}$ be the first meeting time of the atoms at 0 and 1, which is distributed as a rate $\phi(1)$ exponential. Before time $T_{0,1}$, we have $\mu_t^\infty(f) = \frac{1}{3}$. After time $T_{0,1}$, with probability $\frac{2}{3}$, the atom at 0 wins and $\mu_t^\infty(f) = 0$. Otherwise, with probability $\frac{1}{3}$, the atom at 1 wins and $\mu_t^\infty(f) = 1$. Thus, the distribution of $\mu_t^\infty(f)$ on $\mathbb{R}$ - supported at $\{0, \frac{1}{3}, 1\}$ - is given by

$$\exp(-\phi(1)t)\delta(\frac{1}{3}) + \frac{1}{3}(1 - \exp(-\phi(1)t))\delta(1) + \frac{2}{3}(1 - \exp(-\phi(1)t))\delta(0),$$

where here $\delta(x)$ is the point mass at x. Most importantly, note that $\mu_t^\infty(f)$ has no support at $\frac{2}{3}$.

Next, consider the martingales $\mu_t^{(n)}(f)$. Along with $T_{0,1}$, write T_{0,x_n} and $T_{x_n,1}$ for first meeting times between the other two pairs of atoms of $\mu^{(n)}$. We will show that for all n and times $t > 0$, the distribution of $\mu_t^{(n)}(f)$ has non-zero (and bounded below) support at $\frac{2}{3}$.

One way the outcome $\mu_t^{(n)}(f) = \frac{2}{3}$ occurs is if first meeting before time t is $T_{0,1}$, with the atom at 1 winning, and then $T_{x_n,1}$ not happening until after t. Using standard properties of independent exponential random variables, the chance of this is given by

$$\begin{aligned}
&\mathbb{P}(T_{0,1} = \min(T_{0,1}, T_{0,x_n}) \leq t \leq T_{x_n,1} \text{ and 1 wins } T_{0,1}) \\
&\quad = \frac{1}{2}\,\mathbb{P}(T_{0,1} = \min(T_{0,1}, T_{0,x_n}))\,\mathbb{P}(\min(T_{0,1}, T_{0,x_n}) \leq t)\,\mathbb{P}(T_{x_n,1} \geq t) \\
&\quad = \frac{1}{2}\frac{\phi(1)}{\phi(1) + \phi(x_n)}\,(1 - \exp(-(\phi(1) + \phi(x_n))t))\exp(-\phi(1 - x_n)t).
\end{aligned}$$

By continuity $\phi(1 - x_n) \to \phi(1)$ and clearly $\exp(-\phi(x_n)t) \leq 1$ and so

$$\liminf_n \mathbb{P}(\mu_t^{(n)}(f) = \frac{2}{3}) \geq \liminf_n \frac{1}{2}\frac{\phi(1)}{\phi(1) + \phi(x_n)}(1 - \exp(-\phi(1)t))\exp(-\phi(1)t).$$

Therefore, if for our sequence x_n we have $\limsup_n \phi(x_n) \neq \infty$, then

$$\liminf_n \mathbb{P}(\mu_t^{(n)}(f) = \frac{2}{3}) > 0,$$

which implies that

$$\mu_t^{(n)}(f) \not\Rightarrow \mu_t^\infty(f)$$

as each of $\mu_t^{(n)}(f), n \geq 1$ and $\mu_t(f)$ have discrete finite support contained within $\{0, \frac{1}{3}, \frac{2}{3}, 1\}$.

Note that this counter-example applies as long as there is any sequence of points $x_n \downarrow 0$ with $\lim_n \phi(x_n) \neq \infty$. Therefore, for the MC process to be Feller continuous the rate function $\phi(x)$ must satisfy

$$\lim_{x \downarrow 0} \phi(x) = \infty.$$

9.2 Finite Support for Positive Times

In Section 4 we showed that if our initial measure $\mu_0 \in P_{\text{cs}}(S)$ then for all positive times $t > 0$, $\mu_t \in P_{fs}(S)$ and the process evolves as the Metric Coalescent.

Our proof, via a comparison to Kingman's Coalescent, relies heavily on bounding the meeting rates of atoms away from zero using compactness. In the case of

$$\inf_{x>0} \phi(x) > 0,$$

the same comparison to Kingman's Coalescent holds.

Proposition 9.1. *If* $\inf_{x>0} \phi(x) > 0$*, then for any* $\mu \in P(S)$*,* $\mu_t \in P_{\text{fs}}(S)$ *for all* $t > 0$*, almost surely.*

Perhaps ideally this would still hold even without such an assumption, however as we'll see that isn't the case. For the rest of this section, we assume that instead

$$\liminf_{x\to\infty} \phi(x) = 0.$$

We'll then give an example of a countably supported initial measure μ, satisfying any moment condition we'd like, which for all positive times $t > 0$ has non-zero probability of not being finitely supported.

9.2.1 An Example

For simplicity we'll work over $S = \mathbb{R}$ with the Euclidean metric and assume that $\phi(x) \to 0$ as $x \to \infty$. We can then construct an initial measure μ dispersed enough so that $K(t)$, defined as the cardinality of the support of the infinite token process, has

$$K(t) = \infty$$

for some positive times $t > 0$.

Let $r_i \geq 0, i \geq 1$ be a – to be specified – decreasing sequence. We claim that we can then select a sequence of points $s_i, i \geq 1$ in $\mathbb{R}$ such that for each i and $j \neq i$, we have

$$\phi(d(s_i, s_j)) \leq r_i.$$

We choose our initial measure μ to have countable support given by these points, i.e.

$$\mu = \sum_{i=1}^{\infty} m_i \delta(s_i),$$

where

$$\sum_{i=1}^{\infty} m_i = 1.$$

We make no assumptions whatsoever on the mass at each point s_i other than $m_i > 0$.

Initially, as all points $s_i \in \operatorname{supp} \mu$ have positive mass, almost surely there is at least one token (actually infinitely many) with initial location s_i. At time $t = 0$ all the tokens at each point s_i merge into the lowest there. Let $l(i)$ be the lowest token at s_i. Now the total rate q_i at which $l(i)$ merges is clearly bounded

by

$$\begin{aligned} q_i &\leq \sum_{j=1}^{\infty} \phi(d(s_i, s_j)) \\ &\leq \sum_{j=1}^{i-1} r_i + \sum_{j=i+1}^{\infty} r_j \\ &\leq (i-1)r_i + \sum_{j=i+1}^{\infty} r_j, \end{aligned}$$

and is in fact much lower since only finitely many (in fact $l(i) - 1$) of these corresponding merges are actually possibly.

Therefore, the total rate q at which the next meeting occurs is bounded by

$$\begin{aligned} q &\leq \sum_{i=1}^{\infty} q_i \\ &\leq \sum_{i=1}^{\infty} \left((i-1)r_i + \sum_{j=i+1}^{\infty} r_j \right) \\ &= \sum_{i=1}^{\infty} 2(i-1)r_i. \end{aligned}$$

Now we can clearly choose r_i so that the total meeting rate q is finite, for instance $r_i = \frac{1}{i^3}$. Therefore, choosing such an r_i, for all times $t > 0$ there is a positive (i.e. non-zero) probability that $K(t) = \infty$. In fact, we have shown something much stronger - that there is a positive probability of **no** meetings occurring before any time t.

Note something else important about this example: it depends only on the support of μ, not in any way on the distribution of mass of μ. Therefore we can choose the mass m_i so that μ satisfies any moment/concentration condition, for instance

$$\mathbb{E} \exp(d(\xi_1, \xi_2)t) < \infty$$

for all t, where $\xi_1, \xi_2 \sim \mu$.

10 Further Directions

In this section we present some possible further directions of research on the Metric Coalescent process.

10.1 Coming Down From Infinity

Consider a fixed metric space (S, d) and rate function ϕ. Our primary interest in the extended MC process is from its connection to the finite MC process, and so its reasonable to consider for which initial measures this connection holds.

In Section 4 we have seen that for compactly supported initial measures $\mu \in P_{\text{cs}(S)}$ that for all $t > 0$, μ_t is finitely supported almost surely. In Section 9.2

we've seen that even for countably supported (but non-compact) measures, this need not be true.

In the world of Λ-Coalescents, a class including Kingman's Coalescent, under some mild assumptions there is a zero-one law stating that a coalescent has either finitely many blocks for all $t > 0$ almost surely, corresponding to

$$\int_0^1 x^{-1}\Lambda(dx) < \infty,$$

or in the opposite case infinitely many blocks for all $t > 0$ almost surely (Theorem 3.6, [6]). Here Λ is the finite measure on $[0, 1]$ generating the coalescent.

In the world of the Metric Coalescent and token process, the situation is a bit more complicated as there is a clear dependence between when partition blocks merge and their location. At any positive time $t > 0$, it seems likely that the location of the remaining blocks are not independently distributed as μ, since there is a clear bias for "far away" blocks to persist. Nevertheless, we conjecture that an analogous zero-one law holds.

Open Problem 1. *Is there an analogous zero-one law for the Metric Coalsecent, i.e., is it true that for any initial measure $\mu \in P(S)$, that one of the following holds:*

1. *Almost surely, for all $t \geq 0$, $\mu_t \notin P_{\text{fs}}(S)$.*
2. *Almost surely, for all $t > 0$, $\mu_t \in P_{\text{fs}}(S)$.*

If such a zero-one law holds, a natural next step is to consider the class of **finite type** measures $P_{\phi\text{-finite}}(S) \subset P(S)$ given by

$$P_{\phi\text{-finite}}(S) = \{\mu\colon \text{ almost surely }, \forall t > 0, \mu_t \in P_{\text{fs}}\}.$$

We have shown that $P_{\text{cs}}(S) \subset P_{\phi\text{-finite}}(S)$, but that not all countably supported measures are finite type.

Open Problem 2. *Characterize the class $P_{\phi\text{-finite}}(S)$ of finite type measures.*

10.2 Time Reversal

A classical result about Kingman's Coalescent is its duality under a time reversal to a conditioned Yule process[5]. Viewing the Metric Coalescent as a generalization of Kingman's model, it's natural then to consider what can be said about the time reversed Metric Coalescent process.

For a compactly supported initial measure $\mu \in P_{\text{cs}}(S)$ we know (by Section 4) - writing T_n for the first time that $\mu_t, t \geq 0$ has n atoms - that $T_n < \infty$ almost surely for all n. Then, the time reversed process

$$\mu_{T_1}, \mu_{T_2}, \dots$$

is a measure-valued branching process in S with initial condition $\mu_{T_1} = \delta(\xi_1)$, i.e. a point mass at a location chosen from μ. Now by 6.1 we know that

$$\mu_{T_n} \to \mu$$

weakly almost surely. This suggests that the time reversed process can be viewed as a branching process in S conditioned to converge to μ.

Open Problem 3. *Give an explicit description of the time reversed Metric Coalescent process.*

For a local description of the time reversal at $t = 0$, consider the MC on $S = [0, 1]^2$ started from an initial measure μ that's absolutely continuous with respect to Lebesgue measure. We can then view the MC process $\mu_t, t \geq 0$ as a point process on $S \times (0, \infty)$ with the point (s, x) representing an atom $x\delta(s)$. Then, analogous to the case of the self-similar $t \to \infty$ asymptotics in the setting of the classical Smoluchowski coagulation equation [4][11], we make the following conjecture.

Open Problem 4. *For some scaling function $\psi(t) \uparrow \infty$ as $t \downarrow 0$, the point process in a shrinking window around some fixed s_0, rescaled by the map*

$$(s_0 + s, x) \to (s\psi(t), x\psi^2(t))$$

converges as $t \downarrow 0$ to a translation invariant Poisson Point process on $\mathbb{R}^2 \times (0, \infty)$ of some intensity $\delta(s_0, x)$ which is given by the solution of a certain equation.

References

[1] *Real and complex analysis.* McGraw-Hill Book Co., New York, third edition, 1987.

[2] Particle representations for measure-valued population models. *Ann. Probab.*, 27(1):166–205, 1999.

[3] David Aldous, Daniel Lanoue, and Justin Salez. The compulsive gambler process. In preparation, 2014.

[4] David J. Aldous. Deterministic and stochastic models for coalescence (aggregation and coagulation): a review of the mean-field theory for probabilists. *Bernoulli*, 5(1):3–48, 1999.

[5] Julien Berestycki and Nathanaël Berestycki. Kingman's coalescent and Brownian motion. *ALEA Lat. Am. J. Probab. Math. Stat.*, 6:239–259, 2009.

[6] Nathanaël Berestycki. Recent progress in coalescent theory. 16:193, 2009.

[7] Patrick Billingsley. *Convergence of probability measures.* Wiley Series in Probability and Statistics: Probability and Statistics. John Wiley & Sons, Inc., New York, second edition, 1999. A Wiley-Interscience Publication.

[8] Emmanuel Boissard and Thibaut Le Gouic. On the mean speed of convergence of empirical and occupation measures in wasserstein distance. *Annales de l'Institut Henri Poincaré, Probabilités et Statistiques*, 50(2):539–563, 05 2014.

[9] D. J. Daley and D. Vere-Jones. *An introduction to the theory of point processes. Vol. II.* Probability and its Applications (New York). Springer, New York, second edition, 2008. General theory and structure.

[10] Rick Durrett. *Probability: theory and examples.* Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, Cambridge, fourth edition, 2010.

[11] Nicolas Fournier and Philippe Laurençot. Existence of self-similar solutions to Smoluchowski's coagulation equation. *Comm. Math. Phys.*, 256(3):589–609, 2005.

[12] Alison L. Gibbs and Francis Edward Su. On choosing and bounding probability metrics. *International Statistical Review*, 70(3):419–435, 2002.

[13] Alexander V. Gnedin. On convergence and extensions of size-biased permutations. *J. Appl. Probab.*, 35(3):642–650, 1998.

[14] Joseph Horowitz and Rajeeva L. Karandikar. Mean rates of convergence of empirical measures in the Wasserstein metric. *J. Comput. Appl. Math.*, 55(3):261–273, 1994.

[15] David Pollard. *Convergence of stochastic processes*. Springer Series in Statistics. Springer-Verlag, New York, 1984.

[16] Yu. V. Prohorov. Convergence of random processes and limit theorems in probability theory. *Teor. Veroyatnost. i Primenen.*, 1:177–238, 1956.

[17] Daniel Revuz and Marc Yor. *Continuous martingales and Brownian motion*, volume 293 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, third edition, 1999.

[18] V. Strassen. The existence of probability measures with given marginals. *Ann. Math. Statist.*, 36:423–439, 1965.

[19] A. M. Vershik. Random metric spaces and universality. *Uspekhi Mat. Nauk*, 59(2(356)):65–104, 2004.